AD-A121 939 DESIGN OF AIR STRIKE PLANNING AIDS: LESSONS LEARNED FROM THE ONR OPERATIO..(U) ANALYTICS INC WILLOW GROVE PA F A GLENN ET AL. 15 MAY 82 1454-B N00014-79-C-0656 1/2

UNCLASSIFIED F/G 5/1 NL

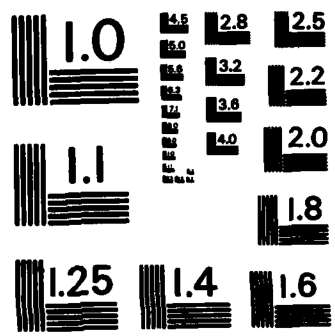

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

2.3 STRIKE PLANNING ORGANIZATION

The organizations most involved in strike planning are:

- Flag or Commanding Officer

Technical Report 1454-B

DESIGN OF AIR STRIKE PLANNING AIDS: LESSONS LEARNED FROM THE ONR OPERATIONAL DECISION AIDS PROGRAM

Submitted to:
Mr. J. R. Simpson
Code 434
Office of Naval Research
800 North Quincy Street
Arlington, Virginia 22217

Contract No. N00014-79-C-0656

15 May 1982

Floyd A. Glenn, III, Ph.D.
Jay M. Bennett
Wayne W. Zachary, Ph.D.

ANALYTICS
2500 MARYLAND ROAD, WILLOW GROVE, PA. 19090

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: 1454-B	2. GOVT ACCESSION NO.: AD-A121939
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle): Design of Air Strike Planning Aids: Lessons Learned from the ONR Operational Decision Aids Program	5. TYPE OF REPORT & PERIOD COVERED: Technical Report 7/16/80 - 4/15/82
	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s): Floyd A. Glenn III, Jay M. Bennett, Wayne W. Zachary	8. CONTRACT OR GRANT NUMBER(s): N00014-79-C-0656
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Analytics, 2500 Maryland Rd., Willow Grove, PA 19090	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS: Engineering Psychology Program, Code 455, Office of Naval Research, Arlington, Virginia 22217	12. REPORT DATE: 15 May 1982
	13. NUMBER OF PAGES: 105
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report): Unclassified
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE: N/A

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

Reproduction in whole or in part is permitted for any purpose of the United States Government.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Decision Aid	Naval Tactics
Air Strike Planning	System Evaluation
Function Allocation	Process Model
Man-Machine interface	Value Model

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report reviews contributions of the ONR Operational Decision Aids (ODA) program to understanding of the decision problems involved in Navy air strike planning and to constructing computer aids for enhancing those decision-making processes. The organization of the Navy strike planning process is summarized. Contributions of individual contractors to the ODA program are reviewed. General observations are offered concerning the technological achievements of the ODA program and the issues in strike planning aid design which require further research.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

ACKNOWLEDGMENTS

The research presented in this report was supported by the Engineering Psychology Programs, Office of Naval Research. The authors wish to acknowledge the direction and important critical contributions to this study made by Mr. J. Randolph Simpson of the Office of Naval Research.

This report contains many judgments concerning the current status of and future prospects for decision-aiding research. These judgments are intended to represent the opinions of the authors and do not necessarily represent those of the Office of Naval Research or of other contractors involved in the Operational Decision Aids program.

TABLE OF CONTENTS

TABLE OF CONTENTS (continued)

TABLE OF CONTENTS (continued)

1. AN OVERVIEW OF THE OPERATIONAL DECISION AIDS PROGRAM

1.1 INTRODUCTION AND BACKGROUND

The Operational Decision Aids (ODA) program of the Office of Naval Research (ONR) was an eight-year effort directed toward the generation of techniques and design concepts for computer-based decision-aiding systems for Naval command and control (C^2) operations. This report presents a review of the accomplishments and lessons learned from the various research efforts undertaken during the eight years of ODA's existence, with particular emphasis on applications for the design of aids for air strike planning.

The ODA program came into existence at the vortex of a number of colliding developments in military systems engineering, computer science, and psychology during the late 1960s and early 1970s. New problems in C^2 decision making arose as a result of advances in sensor and information processing technology and systems engineering. New technologies for sensing, recording, processing, and disseminating information were being implemented in Naval systems and creating a sudden explosion of information for C^2 decision making. Moreover, the general and rapid improvements in weapon system speed and reliability were increasing the pace of events and heightening the consequences of bad decisions. Thus more information was available for consideration in C^2 decisions, but there was less time and less room for error.

Unrelated to these military system events, new developments in computer science have made interactive computing environments efficient and reliable for the first time. Simultaneously, engineering problems in real-time graphics, color graphics, and interactive graphics were being solved, and the graphics technologies were finding a host of applications in the new interactive computing environments. Together the combination of interactive systems and graphics systems opened the door to kinds of applications systems that were impossible only a few years before.

Important developments were also being made in psychology in this period on the understanding and engineering of decision-making behavior. The work of Newell, Shaw, and Simon (1963) brought a new understandability to the inner workings of the complex phenomena of human decision making, and the publication by Raiffa in 1968 of his immensely popular and influential Harvard lectures on decision analysis brought a large audience in contact for the first time with this powerful new normative theory of decision making.

Thus while sensor and information technology provided a new and more complex decision-making environment for C^2, the new developments in psychology (which were being collectively referred to as decision science) and computer science seemed to hold out the promise of a novel solution -- the use of computers to automate the normative decision-science theories and aid human C^2 decision making. The ground-breaking work of Miller, Kaplan, and Edwards (1967) in creating and testing JUDGE, the earliest widely known "decision aid," provided further impetus for research in this area by demonstrating the viability of the approach. It was in this environment that the Office of Naval Research and the Defense Advanced Research Projects Agency (DARPA) began programs in the early 1970s to study the application of normative theories to human decision making in more detail. In 1974, ONR took the additional step of initiating the ODA program to undertake the development of decision-aiding technology and decision-aiding concepts.

1.2 ODA PROJECT ORGANIZATION AND HISTORY

Reflecting the diversity of interests which gave rise to the ODA program, a number of program areas within ONR were brought together to oversee it. The primary ONR groups involved were Engineering Psychology, Information Sciences, Operations Research, and Naval Analysis; Organizational Effectiveness was also involved at the outset. A steering committee composed of individuals from all these groups was formed, and a diverse team of contractors was assembled representing disciplines of:

- decision science,
- human factors engineering,

- computer science (including artificial intelligence),
- operations analysis and operations research, and
- command and control analysis.

The initial goals of the program were established as:

- "A set of command and control decision-aiding methods, strategies, and function assignments, embodied in operational procedures, computer programs, and display formats;
- Quantitative performance criteria and measurement methods for individual and team readiness evaluation, applicable to the establishment of training procedures."*

Thus it was initially felt that decision aiding would be directly applicable to the improvement of C^2 decision making and to the training of C^2 decision makers. The original decision-making domain on which the program focussed was task force command and control. Although other aspects of the program changed (including the level of command decision making addressed), this emphasis on task force C^2 remained. The primary decisions considered were not "heat-of-battle" decisions as much as tactical planning decisions. These are of the kind made before engagements, but which strongly affect the tactical results.

The initial efforts in ODA were directed toward identifying a representative set of task force C^2 decisions around which further aid development could proceed. A scenario was developed (Payne and Rowney, 1975) which considered a task force operation focussed on an island named ONRODA; a wide variety of options were available to the task force commander in this scenario. Individual decisions which might be required in the course of this scenario (or subtle variations of it) were then identified and investigated for aid development by various contractors. The original "ONRODA" scenario resembled in many ways a

*Commerce Business Daily, Issue PSA-5916, September 29, 1973.

specialized Mediterranean situation but proved to have broad applicability, at least in terms of generating decision-aiding situations. It is interesting to note the similarity of the ONRODA scenario to the current situation in the Falkland Islands, where many of the decisions identified and pursued in ODA for aid development are, in fact, being faced by the commanders of the British fleet.

In 1978, a secondary program was begun by ONR to study human information processing and decision making in specific C^2 contexts, with the objective of developing decision aids capable (perhaps with minor modifications) of operational use. Among the specific C^2 problems studied in the secondary program are:

- air antisubmarine warfare,
- attack submarine command,
- marine amphibious assault brigade tactical operations, and
- task force antiair warfare.

This more recent program differs from ODA in that it focusses on the development of fleet-usable aids and the development of techniques for the design of fleet-usable aids. On the other hand, ODA was more general and explored the generation of decision-aid techniques and methodologies without the constraint of producing fleet-acceptable products. Throughout this report, certain of the research efforts produced in this "spinoff" program are cited and discussed along with the ODA products.

1.3 ORGANIZATION OF THIS REPORT

For the most part, the decisions identified within the ONRODA scenario involved aspects of carrier-based air strike planning, and thus most of the ODA aids attempt to improve some decision within the task force air strike planning process. To provide a focus for the more detailed review of ODA, this area of strike planning is used in this report as the thread by which many of the ODA efforts are tied together.

Section 2 provides an in-depth review of what has been learned about strike planning decisions and the overall strike planning process in the course of the ODA program. Section 3 then presents a detailed review of all ODA projects which are in some way related to carrier-based, air strike planning. While the decision aids and their evaluations described in Section 3 represent an important result of the ODA program, they are not the only result or even perhaps the primary result. Through all the efforts to develop specific decision aids, a great deal of technology for aiding C^2 decisions has been developed and a great deal about the process of decision-aid design has been learned. These kinds of ODA program products are discussed in Section 4 of this report.

Another of the results of the ODA program, though perhaps an unintended one, has been the structuring of the discipline of decision aiding. The directions taken by ODA have largely defined the future direction of decision-aiding research. Even though ODA has covered a great deal of ground in the last eight years, it has in many ways only scratched the surface of this large domain. Many issues remain unresolved, and much work remains to be done. Section 5 considers some of the most pressing questions which still remain to be answered after ODA; for the most part, these are questions which arose from research undertaken during the ODA program. Section 6 then presents a summary of this review and some conclusions.

2. OVERVIEW OF AIR STRIKE PLANNING

The air strike planning process encompasses a broad range of decision-making subproblems performed by several levels of command. The goal of the process is a strike plan which:

- Is feasible within time and resource constraints
- Maximizes damage to enemy capability
- Minimizes destruction of friendly resources

Since the maximization of enemy damage and the minimization of friendly losses can be conflicting objectives, one of the critical decisions in the process is the evaluation of trade-offs between disparate types of strike outcomes such as own and enemy force attrition and other tactical situation factors.

2.1 AIR STRIKE MODES

Air strikes operate basically in one of two different modes:

- Force Projection -- Air strikes may be used to aid land operations through close support of troops. Their goal may also be the destruction of enemy land installations such as missile sites, air bases, surveillance sites, and supply centers. Bridges and roads may be damaged to hinder enemy movement of troops and supplies. The objective may also be simply the destruction of grounded airplanes or docked ships.

- Sea Control -- Air strikes are employed to damage and destroy enemy naval vessels to establish control over the blue water environment. However, the destruction of the vessels may be only secondary to another objective (e.g., forcing the retreat of an invasion force).

Each of these modes entails different problems and decisions. For example, terrain is an important factor in force projection air strikes while this factor plays no role in blue water environments. Within each mode, weighing the expected results against the possible cost is difficult. It is hard enough to balance losses of friendly planes against losses of enemy planes, ships, or troops, but the problem becomes even greater when such friendly losses are balanced against the destruction of a bridge or the forced retreat of enemy troops.

2.2 STRIKE PLANNING SITUATIONS

Strike planning is influenced by the contextual situation in which the planners are placed. The time given for plan development and the uniqueness of the combat situation are influential factors. Some plans are developed for a specific current situation. An example might be the destruction of an enemy task force threatening the friendly task force. Other plans are developed for repetitive operations such as the repeated bombing of roads acting as supply conduits. Developing a totally new plan in unique circumstances is much more difficult than modifying a standard plan used repetitively. Contingency plans are often developed for general situations which may occur in the near future so that they may be easily adopted when the appropriate situation arises. Contingency planning places less pressure on the planners by relaxing the time constraints on the period allocated for plan development. Estimates of time needed to currently plan a strike range from two to six hours. While contingency plans may be developed many days before they are required, most strike planning development must be accomplished in a matter of hours. Given that an air strike plan has been developed, the planners may also want to develop a secondary backup plan in case an event or condition arises which destroys the foundations of the main plan. The planners may feel a backup plan is necessary when uncertainty about a key factor such as weather is large.

2.3 STRIKE PLANNING ORGANIZATION

The organizations most involved in strike planning are:

- Flag or Commanding Officer
- Operations
  - -- Strike Operations
  - -- CV Intelligence Center (CVIC)
  - -- Combat Information Center (CIC)
- Air Operations
- Air Wing

In addition, representatives of these organizations and others are combined into two groups for the special purpose of strike planning, the Strike Planning Board and the Strike Planning Team. The titles and compositions of these groups are not standard. Although the strike planning processes are basically the same throughout the Navy, differences between fleets (and even between ships) arise in allocation of responsibility and in the involvement of departments and their representatives. As an example, the determination of resources required for force defense is performed differently in the Atlantic and Pacific fleets. This task is performed only at the flag level in the Pacific whereas in the Atlantic the Strike Planning Board is responsible subject to approval at the flag level. The responsibility for future and current operations is another case in point. Strike Operations is the organization generally allocated both operations in the Pacific fleets whereas Air Operations assumes responsibility for current operations in the Atlantic fleet. Since no authority in the Navy has made a definition of the strike planning procedure, the process has been subject to the vagaries of evolution within the fleets.

2.4 AIR STRIKE PLANNING SYSTEMS

Currently, there exists no single system to support all the various tasks involved in air strike planning. Most planning is performed manually

with support from programs operating on programmable calculators. Some of the applicable programs in the PROCAL library include:

- Joint Munitions Effectiveness Manual (JMEM) Basic Analysis -- weapon selection aid
- A/C NAV -- routing and fuel use calculation aid
- LOADPLAN -- ordnance loading aid

The Naval Intelligence Processing System (NIPS) includes a data base to provide intelligence information (such as enemy order of battle) for strike planners. The Logical Information and Display of Aircraft (LINDA) system is a prototype system developed for retrieval of information on aircraft availability; LINDA has not yet been accepted for use on all carriers. The systems in place do not interact with one another so that transfers of data from one to another are time-consuming and prone to error.

2.5 STRIKE PLANNING PROCESS

The air strike planning process is initiated by the Flag or Commanding Officer of the task force who outlines the following task requirements:

1. Targets (location and identity)
2. Degree of destruction required for each target (expected level of destruction or probability of achieving a specified level of destruction)
3. Rules of engagement
4. Time over target
5. Available forces
6. Command and control constraints
7. Reconnaissance required

The Commanding Officer (CO) passes these requirements to a Strike Planning Board composed of representatives from Strike Operations, Weapons, Air Operations, Intelligence, Navigation, Combat Information Center, and Flag Officer. The first major planning decision is whether the mission will consist of one large strike or several strikes in cyclic operation. A generalized strike schedule is then developed to provide a working basis for the rest of the planning.

The Strike Planning Board performs a detailed analysis of the strike's goals in light of the tactical situation. Terrain, target characteristics, enemy order of battle, wing availability, weapon availability, and force defense requirements are all examined to determine the attainability of the goals set by the CO within the constraints also set by the CO. Any problems may be discussed with the CO to relax the constraints and/or lower the goals.

If the goals appear feasible, the next major objective of the Strike Planning Board is to determine the weapons needed to achieve the required degree of destruction upon each target. This decision is crucial to the mission and influences many other decisions. Using the Joint Munitions Effectiveness Manuals (JMEMs), target destruction may be estimated when using specified weapons delivered in a specified mode by a specified platform. Many problems may be encountered:

- The best weapons for use may be unavailable or in limited supply.
- The best delivery mode may involve too great a risk for the attack aircraft because of enemy defenses or the nature of the terrain.
- The platforms (or special equipment) needed to deliver the weapons may be unavailable.
- The extra weight and drag of the weapons may strain fuel limitations.

Any of these problems may prevent the use of the "best" weapons (i.e., those having the greatest probability of achieving desired target destruction). Generally, a "compromise" weapons mix is selected to achieve target destruction within constraints such as those mentioned above. Unfortunately, the structure of the JMEMs does not facilitate the search since they are designed only to guide the evaluation of specified weapons mixes and do not lead the planner to a weapons mix appropriate for the strike goals.

Once the Strike Planning Board finds a feasible weapons mix achieving the strike objectives, the remainder of the detailed planning is performed by a Strike Planning Team composed of the Strike Leader, Missile Suppression Leader, Navigation Leader, EW Specialist, representatives from each aircraft squadron, the Air Intelligence Officer, and any additional experts necessary to plan the mission. It is the Strike Planning Team's job to complete all the details of the plan outlined by the Strike Planning Board.

A check is made of the availability of weapons specified by the Strike Planning Board. These weapons are assigned to aircraft for delivery to their respective targets. A check is made that the weapons mix assigned to an attack aircraft has a feasible configuration in light of the availability of special installation equipment such as mounting racks. If the initial weapons mix is not feasible, new selections of weapons are made and JMEM calculations are performed until a mix is found which optimizes target destruction within availability limits. The team details the penetration route to the target and the escape routes back to the task force. Here the planner has the difficult task of designing routes which provide the least exposure to sensors and defenders without consuming excessive fuel. Deception in route planning is important in increasing the air strike's chances of success. The target area is examined to determine the best approach to maximize the attack's effect and to evade enemy defenses. A delivery mode must be selected which permits the greatest probability of target destruction without posing too great a risk to friendly forces (within restrictions imposed by weather and terrain). The tactics used in attacking the targets are made with reference to TACNOTES as guidelines and

to the experience levels of the personnel involved. The team also configures the mission, determining especially whether the strike is performed in a single wave or in multiple waves.

With aircraft routes and weapon loads as input, the mission fuel requirements are calculated with the use of NATOPS manuals. If more fuel is needed than the aircraft can normally carry, three solutions are available to the planners. First, auxiliary tanks may be carried by the aircraft; in this case, the weapon load of each aircraft is rechecked to determine if the aircraft can carry its assigned weapons as well as the auxiliary tanks. Second, a requirement for tanker aircraft to rendezvous with returning aircraft low on fuel may be specified. Third, if neither of the preceding solutions is feasible, the routes and/or weapon loads must be changed.

In addition to tanker requirements, the Strike Leader determines the necessity for support in two other areas: Search and Rescue (SAR) and Airborne Early Warning (AEW). If SAR support is deemed necessary, helicopters, ships, and Rescue CAP (RESCAP) aircraft are assigned and coordinated with respect to time and likely recovery locations. Terrain and weather are important considerations in planning SAR support. The Airborne Early Warning (AEW) requirement is typically satisfied by dedicating one E-2C aircraft to support the strike mission by providing electronic warning and control. Since AEW aircraft are highly vulnerable to attack, careful consideration is given to their placement to maximize their support capabilities without excessive risk. Upon examination of data on enemy defenses, tactical and electronic countermeasures are selected. Emission Control (EMCON) may be enforced to prevent enemy detection of the strike force. Fighters are assigned as escorts for the attack aircraft to counter the threat of engagement by enemy air interceptors; their use in the strike must be balanced against their value in task force defense. The assignment of aircraft for jamming enemy communications is considered as are tactics to counter jamming by the enemy.

Once the Strike Planning Team has completed the plan, the Strike Planning Board reviews it. Upon passing this check, the Officer in Tactical

Command (OTC) is briefed and, if satisfied, gives final approval to the plan's implementation.

2.6 PROBLEM AREAS OF AIR STRIKE PLANNING

As Section 2.5 indicates, strike planning involves a multitude of decisions each affecting several others and ultimately affecting the probability that the strike's goals will be achieved. Some decisions such as penetration routing are hard to make because of the difficulty in estimating the differences in alternatives with respect to their effect on the mission's probability of success (e.g., how much does route A improve the chances of success over route B, if at all?). Other decisions such as delivery mode selection involve balancing units of gain against non-comparable units of loss (e.g., is delivery mode A better than delivery mode B if it increases the probability of success by p percent at the expected expense of x more aircraft lost?). Still others such as weapon loading require satisfying several complex constraints in addition to optimization of damaging effect (e.g., what weapons mix is best given the aircraft available, the special equipment available, the delivery mode, and the fuel requirements?). To make some decisions, time-consuming calculations which are prone to error must be made (e.g., fuel consumption calculations in route planning and target destruction in weaponeering). All of these factors combine to make air strike planning an extremely difficult task.

The intricacy of the problem is reflected in the hierarchy of command levels (Flag, Strike Planning Board, Strike Planning Team) required to create the strike plan. Each level attacks the problem in greater detail and double-checks the work of the higher level which passed on the problem; this process is climaxed by the final review where the top level (Flag) checks the detailed plan produced by the lowest level (Strike Planning Team). This process consumes large numbers of man-hours of high-level personnel often in situations where their time is critical.

Given the constraints on time, the process tends to produce plans which are feasible and satisfactory rather than being optimal. Optimization

of air strike plans necessitates experimentation with ideas and hypotheses for each subproblem. Optimal solutions within the subproblems must then be tested with solutions obtained for other subproblems. Several iterations through the entire process may be needed to bring an entire plan close to optimality.

Thus, the current air strike planning process demonstrates difficulties in:

- Consuming large amounts of clock time
- Using great quantities of man-hours of valuable personnel at critical times
- Optimizing plans
- Non-standardization of procedures through the Navy
- Inefficient data-gathering support
- Lack of quantitative means to evaluate tactics and strategy
- Comparing expected results from competing alternative plans

3. SUMMARY OF ODA PRODUCTS FOR STRIKE PLANNING

Although support for the strike planning process was not an explicit objective of the ODA program, most of the ODA efforts produced tools or concepts potentially applicable to aids for strike planning. Several projects developed aids which specifically addressed problems of strike planning (i.e., projects conducted by Analytics, SRI International, and Integrated Sciences Corporation). Other projects developed general aids applicable to a broad range of tactical problems which could include aspects of strike planning (i.e., projects conducted by SRI International; Decisions and Designs, Inc. (DDI); the Wharton School of the University of Pennsylvania; Grumman Aerospace Corporation; and CTEC, Inc.). One project (initiated by General Research Corporation and later transferred to Decision-Science Applications) undertook the development of an electronic warfare decision aid which addresses a problem that is distinct from strike planning but which offers several techniques that could be applied to the strike planning domain. Three additional projects (those conducted by Applied Psychological Services, the Navy Personnel Research and Development Center (NPRDC), and Analytics) addressed the problem of evaluating the usefulness of decision aids, with two of the three evaluations that were performed being specifically concerned with strike planning issues.

Brief overview descriptions of individual ODA projects are offered here to provide a context for the subsequent discussion of specific component products that are relevant to strike planning. Projects will be distinguished by contracting organization, and distinct research efforts within each organization will be identified as much as possible. It should be recognized, however, that considerable redirection of effort occurred in several projects as promising lines of development were discovered, so that a number of distinct efforts were occasionally pursued under a single contract by a given contractor.

This review is not intended to cover every issue and technique investigated in the ODA program, but rather to offer a general picture of the major thrusts of each project that were or could be related to strike planning problems.

3.1 ANALYTICS

Analytics undertook three ODA projects -- development of general techniques for measuring the effectiveness of decision aids, development of an aid for selecting the time to launch an air strike, and development of an aid for air strike planning. The present report represents one component of the latter effort.

The measures-of-effectiveness project (Martin, 1976) surveyed the general problem that must be addressed in evaluating decision aids and then focussed on development of a technique for evaluation of one particular but very important class of decision aids -- aids for situation assessment. It was determined that evaluation of such aids requires an assessment of the degree to which the aid fosters the generation of useful, discriminating information regarding the characteristics of a tactical situation. A general mathematical formula based on information theoretic concepts was described for measurement of aid effectiveness. However, practical empirical methods for implementing this technique were not developed, and the technique was not employed in any subsequent evaluations of ODA products.

The strike timing aid project (Epstein et al., 1977) constituted an attempt to develop an alternative approach to the type of decision analytic aid which DDI was developing. The DDI aid had been presented as a general technique for integrating knowledge about utilities and probabilities of possible states of the tactical situation and possible decision-maker actions in order to determine a preferred alternative. (The DDI project will subsequently be described in detail.) The principal exemplary application for the DDI aid in DDI reports on the project was to the problem of determining whether a Blue force commander should launch an attack against Red now, prepare to launch an attack in a short time, or make no immediate plans to attack based on an assessment of whether Red

was planning to attack Blue, planning to harass Blue forces, or continuing routine operations. Rather than treat states of the world and action options as discrete alternatives for which probabilities and utilities must be separately generated, Analytics suggested that the decision space be treated as continuous, with the variables of interest being related by mathematical models for the processes of concern. The problem addressed by DDI was transformed into one in which the problem for the Blue force commander was to select a time for launching a strike against Red based on assessments of Blue force readiness, Red force readiness, and predicted weather conditions. Utility weights were assigned to Red and Blue force components, and readiness estimates were represented as probability distributions for possible numbers of available units for each type of force component. An engagement model was developed to describe attrition of force units as a function of numbers and capabilities of engaging force units and prevailing weather conditions. The model was exercised in Monte Carlo fashion as a computer simulation, and results were presented both in terms of specific predicted force unit attrition and as summary utility values. Options for the user to perform sensitivity analyses on major variables and to examine results at different stages of the engagement process were offered. Color-coded bar graph displays indicating both central tendency and uncertainty (i.e., variability) of outputs were offered along with tabular displays of the same data. The aid was called the Air Strike Timing Decision Aid (ASTDA).

The project to develop an integrated aid for strike planning (Glenn and Zachary, 1979; Glenn and Bennett, 1980) grew from recognition of the fact that many of the ODA efforts were producing aids and aiding components that specifically addressed aspects of the strike planning process. The initial effort focussed on the integration of three aids that were directly concerned with strike planning -- ASTDA, the Strike Outcome Calculator of SRI International, and the Route Planning Aid of Integrated Sciences Corporation. It quickly became clear, however, that an integrated aid could benefit substantially from incorporation of components of many other ODA products. It also became evident early in this effort that it would be necessary to draw some aiding components from outside of the ODA program in order to produce a truly comprehensive aid for

any significant domain of strike planning (e.g., components would be required for weaponeering, aircraft fuel calculations, and cartographic data management). Separate aiding concepts were conceived for addressing the problems involved in planning tactical details for individual strikes (Glenn and Bennett, 1980) and more general tactics for protracted multi-strike campaigns (Glenn and Zachary, 1979). A cursory effort was made to assess operational fleet requirements for these types of aids, and it was determined that these requirements are quite complex, so that considerable further investigation of fleet needs should be undertaken before implementation of a comprehensive strike planning aid is pursued.

3.2 APPLIED PSYCHOLOGICAL SERVICES

Applied Psychological Services (APS) conducted two projects to evaluate aids produced in the ODA program. These projects served several purposes simultaneously -- to provide estimates for candidate user organizations of the potential usefulness of each aid, to provide feedback to aid developers concerning the value of component features of the aids, and to develop effective methodologies for evaluation of decision aids. The two aids that were evaluated were the Air Strike Timing Decision Aid (ASTDA) developed by Analytics and the Electronic Warfare (EWAR) decision aid developed by Decision-Science Applications.

The evaluation of ASTDA (Siegel and Madden, 1980) dealt with three factors -- aid components available to the user, relevant decision-making experience of the user, and problem difficulty. Problems were presented to experimental subjects in terms of predicted force readiness and weather conditions at each of six possible strike launch times, and the task was to determine a preferred launch time (Glenn, 1978). Different aiding conditions were employed so that subjects would use either no outputs of the aid, force attrition predictions only, expected utility predictions only, or all aid predictions. Availability of uncertainty (i.e., variability of output variable distributions) was also manipulated. Decisions were evaluated with respect both to the optimum defined by the engagement and value models in ASTDA and to the

consensus of a panel of experienced Naval officers. It was determined that the aid tended to enhance decision making significantly by both criteria. It was observed, however, that experienced Navy decision makers, in particular, felt that the utility model incorporated in the aid was not helpful and that only predicted outcomes should be offered in an operational system.

The evaluation of EWAR (Madden and Siegel, 1980) employed a design similar to that used for ASTDA, with three factors being addressed -- aid availability (all or none), training offered to subjects, and problem difficulty. As will be described subsequently in greater detail, EWAR aids a tactical analyst in developing a radar emission control (EMCON) plan for a task force that achieves a desirable trade-off between surveillance effectiveness and information denied to the enemy. Problems consisted of specification of task force composition, values of individual ships, radar resources available and their detection capabilities, and probabilities of specific types of enemy attack. The subject's task in each case was to determine an electronic order of battle specifying which radars to turn on and which to turn off. Decision effectiveness was measured in terms of the fraction of the task force that survived a simulated enemy attack. It was determined that plans developed with the aid were significantly superior to plans developed without the aid and that the trade-off analysis feature of the aid was particularly useful.

3.3 CTEC, INC.

CTEC, Inc., was tasked to provide an experimental data base for the ODA program to be used in conjunction with the test bed implemented at the Wharton School. The data base (CTEC, 1976) was designed to support the evaluation of decision aids in the context of scenarios provided by SRI International. The data base constituted a sanitized variation on current Navy data bases, emulating both their structure and content. Although this data base was used very little as a resource for decision-aiding software, it provided important guidance to decision-aid developers concerning the characteristics of the informational environment in which Navy tactical decision aids must be implemented.

In order to provide a full spectrum of support to the decision aids of the ODA program, CTEC also produced a collection of mathematical models for generic, tactically significant processes (CTEC, 1977). These include models for such processes as military combat engagements and radar detections of aircraft. Variants of these models are incorporated in several of the ODA decision aids. The idea that generic process models should be closely associated with and perhaps integrated into a general Navy tactical data base represents recognition of the need for use of common models by related decision-aiding products. Use of a common modeling base is important for focussing of validation and calibration efforts, facilitation of interfacing between aids, and minimization of data requirements.

3.4 DECISIONS AND DESIGNS, INC. (DDI)

DDI undertook the broad task of conceiving and implementing decision-aiding tools based on the techniques of decision analysis. They explicitly addressed the problems faced by the Navy task force commander and his staff (Brown et al., 1974). As a context for consideration of aiding requirements, four representative scenarios were identified, two of which concerned air strike planning. A variety of decision analytic techniques were proposed as candidate aids for making the kinds of tactical decisions represented in the scenarios. Basic candidate aids were described for value assessment, probability assessment, construction and analysis of decision trees, and implementation of decision rules. Several aids which integrated two or more features of these basic aids were also developed. Aids were described for both mission planning and mission execution roles (Brown et al., 1975). All aids were conceived as applications of general decision-making techniques, with a consequent broad range of applicability being intended. Illustrative applications of each aid to realistic decision problems of a task force commander were constructed.

A large proportion of DDI's aid development effort was devoted to one composite decision aid called TACAID, which constitutes an integration of probability modeling, value modeling, and decision rule techniques to address a broad class of mission execution decisions (Peterson et al., 1976; Peterson et al.,

1977; Barclay et al., 1979). TACAID employs a straightforward implementation of multi-attribute utility theory in the form of a linear, additive value model and a probability estimation process based on the Bayesian updating technique. An innovative display technique is used to present summary decision analytic information to the user. The displays consist of triangles which represent the probability state space for three exhaustive, independent, and mutually exclusive conditions which significantly influence the decision to be made. (The stipulated restrictions on the three-dimensional state space imply that the probabilities of the three independent conditions must sum to one so that the range of feasible probability triads is limited to a triangular portion of a two-dimensional plane.) The utility model for action options, as a function of which of the three tactical conditions prevails, serves to partition the triangular region so that location of the probability triad (i.e., a point in the probability state space) in a given region indicates that a particular action option yields the maximum expected utility and, hence, is optimal. The point in the triangle which represents the current probabilities of the three possible conditions can be observed to move as indicator information is used, via Bayesian updating, to determine a new probability triad. Thus the user can observe the sensitivity of the optimal action selection to unfolding indicator data or other movement in the probability space (e g., to analyze regions of uncertainty around current probability estimates).

The principal illustrative decision problem to which DDI applied TACAID was a situation in which the task force commander must decide what offensive action to implement in response to indicators pertaining to the enemy's intent to attack the task force. The Bayesian updating technique was used to estimate probabilities of hypothetical events characterizing enemy intentions (this technique is called Bayesian hierarchical inference). Enemy intentions are assumed to be in one of three exhaustive and mutually exclusive states -- continuing routine operations, intending to harass, or intending to attack. Possible action responses that are considered by the commander are similar to the options assumed for the enemy -- routine operations can be continued, preparations can be made to launch an attack at a later time, or an attack can be launched immediately. TACAID was evaluated for this application, informally by DDI and more

formally by the Navy Personnel Research and Development Center (NPRDC). These evaluations disclosed two serious problems with the Bayesian updating component of the aid relating to aid assumptions concerning the independence of indicators and the stationarity of the situation being assessed. The nature of these problems will be discussed in detail in the subsequent review of the NPRDC evaluation effort.

It was recommended in the final DDI contribution to the ODA program (Barclay et al., 1979) that TACAID be expanded to treat more than three situational conditions and that techniques be developed and incorporated to allow for dependent indicators and non-stationarity of the tactical situation. Specific approaches were suggested for these further developments, and it was noted that the expanded aid would be much more complicated than the initial TACAID and would require substantial user training.

3.5 DECISION-SCIENCE APPLICATIONS, INC./GENERAL RESEARCH CORPORATION

A team of investigators initially affiliated with General Research Corporation and later with Decision-Science Applications, Inc. (DSA), conducted two distinct projects for the ODA program. The first effort surveyed the potential applications of mathematical techniques to Naval tactical decision-aiding. The second effort developed a decision aid to help in formulation of radar emission control (EMCON) plans for Naval task forces.

The project to investigate mathematical decision aids (Pugh, 1976) identified two general types of mathematical applications for decision-aiding corresponding to analysis of outcome models for military situations and computational assistance in implementing the methods of decision analysis. It was determined that outcome models (i.e., models which predict outcomes of processes which are tactically significant) can be applied only to problems which are highly structured and at least somewhat routine. Promising areas for development of outcome calculators include fleet air defense planning, air strike planning, and antisubmarine warfare. The methods of decision analysis which provide general guidance in characterizing and analyzing decision problems are particularly applicable to unique, high-level problems.

In the general investigation of mathematical decision aids, it was determined that a decision aid for fleet air defense planning would be both useful and technologically feasible, so a subsequent effort addressed the development of such an aid (Pugh et al., 1977; Densmore et al., 1978; Noble, 1980). The fleet defense aid focussed on the problem of aiding a tactical analyst to construct a radar EMCON plan for a task force. The aid, called the Electronic Warfare decision aid (EWAR), assists in the management of plan data, in the assessment of plan consequences in terms of surveillance effectiveness and information about the task force disclosed to the enemy, and in the analysis of the trade-off between surveillance effectiveness and information given away. Candidate plans are analyzed first to determine the degree to which the enemy will be able to identify the high-valued units of the task force and direct their attacks against them. A simulation model is then used to generate enemy strikes, to determine when they are detected by task force radars, and to predict resulting attrition of the task force. By assigning values to individual task force units, total values lost to the task force in each strike can be determined and an overall value of each EMCON plan can be calculated. A variety of separate scores are determined for surveillance effectiveness and information denial criteria, and a trade-off analysis display is offered to help the analyst identify an appropriate compromise between the extreme planning options. An assortment of graphic and tabular displays provides for user access to aid inputs, outputs, intermediate results, and underlying assumptions. An innovative technique based on graphic representations called influence diagrams has been developed to enhance the user interface by efficiently displaying an overview of the structure, capabilities, and limitations of the aid (Noble and Pugh, 1981).

3.6 GRUMMAN AEROSPACE CORPORATION

The Grumman Aerospace Corporation undertook two decision-aid development efforts in the ODA program, both of which are potentially relevant to a strike planning aid. The first consisted of the construction and evaluation of a decision aid called the Options Selection Checklist. The second consisted of the conceptual development of geographical displays of tactical information relevant to routing decisions and selection of an operations area for a task force.

The Options Selection Checklist is a computer-based mechanism for implementing a multi-attribute utility analysis of a decision problem (Kalenty and Lockwood, 1976; Kalenty et al., 1977). The aid offers a "canned" set of factors (or attributes) that the decision maker should consider for each particular type of decision problem. For each option available to the decision maker, figures of merit (or utilities) must be supplied to represent estimated values of the option relative to each factor. The decision maker must then provide weight values for the factors, which are used to calculate a composite figure of merit for each option. The composite figure of merit can then be used by the decision maker as the basis for selection of a preferred option.

The Options Selection Checklist was evaluated in a variant of the ONRODA scenario in which the Blue task force commander must select one ship from the task force to accomplish a "tattletale" mission (Kalenty et al., 1977). The tattletale ship would be required to follow and observe an enemy (Red) ship, taking action to neutralize any activity which would threaten the task force. Three decision-aiding variations were offered to experimental subjects:

- An unaided condition in which the subject had access only to a variety of displays representing resources and capabilities of the Blue and Red forces (these displays were designed to simulate the structure and content of displays offered by current Navy command and control systems).
- An aided condition in which the subject was required to supply component figures of merit and weight values for determination of a composite figure of merit.
- An aided condition in which figures of merit and factor weights were provided to the subject for review and revision before the composite figure of merit was determined.

All of the data displays offered in the unaided condition were also available in the two aided conditions. Subject performance in the unaided condition was judged to be quick but erratic, with important factors frequently being ignored. Subjects expressed a preference for some combination of the two aided conditions in which objectively determinable figures of merit (e.g.,

radar surveillance capability of a ship) would be calculated by computer algorithm but more subjective figures of merit (e.g., capability of a ship to provide task force defense) would be left to the decision maker. There was also general agreement among subjects (who were all experienced in Navy tactical operations) that compelling evidence of aid validity in terms of both reliability and accuracy of data inputs and resulting recommendations would be required before operational use of the aid could be expected.

In a second effort, Grumman developed concepts for a variety of geographical displays that could support decisions concerning selection of an operations area for a task force and a transit route of task force ships (Kalenty et al., 1977). Display concepts were defined for summarizing own and enemy force capabilities for weapon delivery and radar surveillance. Composite displays were also described to represent figures of merit for the total threat situation and the combined factors of task force vulnerability and effectiveness. These composite displays used a varying density of dots on the geographic display to code figure-of-merit values.

3.7 INTEGRATED SCIENCES CORPORATION (ISC)

Integrated Sciences Corporation (ISC) conducted two related project efforts for the ODA program. One protracted effort investigated the applicability of a variety of decision-aiding techniques to the determination of an optimal action selection in a complex multi-dimensional space. The second effort addressed the implementation of techniques developed in the first effort as a decision aid for air strike planning.

The ISC investigation of optimization techniques focussed on the problem of finding an optimal route for an aircraft in flying from a launch point to a target area through a field of enemy sensors. Optimization was to be performed relative to a complex utility function which incorporated the factors of fuel consumption and probability of detection by the enemy. For most of the investigation, it was assumed that the aircraft route would be defined as a sequence of connected straight-line segments and that the aircraft would fly at

a constant speed along each segment but could change speed in moving from one segment to the next. While this optimization problem clearly represents the kind of routing problem that must be addressed in air strike planning, it was intended from the outset that it would be a simplification of the operational problem which could be efficiently used for development of and experimentation with candidate decision-aiding concepts.

A series of experimental issues were sequentially investigated by ISC in order to determine appropriate features of a route planning aid. Principal issues that were examined include the following:

(1) What is the best technique for constructing a composite detection rate map for a given geographical distribution of radar sites? An experiment determined that humans can perform very well in using an interactive computer graphics system to construct composite detection rate contours from displayed detection rate contours for individual sensors (Irving et al., 1977). However, a computer algorithm was developed which could accurately determine composite detection rate contours so that it would not be necessary for subsequently developed route planning aids to depend on human input of composite contours.

(2) Could human route planning performance with an interactive computer graphics system be improved though use of a mathematical optimization algorithm? It was experimentally determined that some improvement could be obtained through the use of a dynamic programming algorithm (Irving et al., 1977). However, the improvement attributable to the algorithm was fairly small, which may have been the result of the simplicity of the experimental problems or other experimental factors.

(3) What is the appropriate allocation of optimization responsibilities between the human and a computer algorithm relative to generating initial candidate routes and then iteratively improving upon them? An experimental study of this issue determined that a non-linear programming algorithm to which a human supplied initial candidate routes produced superior performance to unaided humans and to alternative, fully automated systems based on non-linear programming and dynamic programming algorithms (Walsh and Schechterman, 1978). The distribution of functions for candidate route generation and iterative improvement between man and machine was termed Operator-Aided Optimization (OAO).

(4) How does performance of the OAO system based on non-linear programming compare to performance of a system in which the human generates initial candidate routes and iteratively improves on

them with the computer only providing feedback on the utility of each route? (This latter aiding option is called Iterative Manual Optimization -- IMO.) Experimental evidence indicated that OAO performance was significantly superior to IMO performance, but that the difference was rather small (Schechterman and Walsh, 1980).

At the conclusion of this series of experiments, the appropriate features for an operational route planning aid were still somewhat indeterminate, but the OAO and IMO options were deemed worthy of further examination in the context of more realistic route planning problems.

The second ISC effort in the ODA program proceeded to conceive and develop an operationally viable strike planning aid based on the route planning concepts investigated in the earlier effort (Walsh et al., 1981). Experienced strike planners were consulted to determine desirable features for such an aid. It was determined that an Air Strike Planning System (ASPS) should incorporate all factors addressed by the earlier route planning research and also aircraft altitude, terrain characteristics, enemy defenses, electronic countermeasures, versatile area mapping, weapon loading, and determination of tolerable corridors about optimized paths. Specific recommendations were offered for hardware requirements for system implementation. A variety of user interface innovations were also described.

3.8 NAVY PERSONNEL RESEARCH AND DEVELOPMENT CENTER (NPRDC)

NPRDC undertook the evaluation of TACAID, developed by DDI in a simulated operational setting (Gettys et al., 1976). The evaluation employed the particular aid application and general scenario that had been posed by DDI in early expositions of the aid -- selection of an offensive posture by a Blue task force commander in response to unfolding indications concerning possible intentions of the enemy (Red) to attack the Blue task force. In addition to evaluating the aid developed by DDI, NPRDC also evaluated a variant of the aid which employed an alternative interface to the user, offering information concerning the degree of risk (predicted variability of outcome values) associated

with action alternatives. Whereas the DDI TACAID provided indication only of which action would yield maximal expected utility, the NPRDC-developed variation also displayed the probability that each feasible utility result would be achieved. Both versions of TACAID employed the same mechanisms for the processing of probabilities and utilities.

The design of the evaluation experiment consisted of three aiding conditions (no aid, the DDI version of TACAID, and the NPRDC version of TACAID) and two risk levels for experimental scenarios (low risk, in which task force responsibilities could be assumed by other Blue task forces, and high risk, in which the Blue task force had no proximal support). Experimental subjects, who were all experienced in Navy tactical decision-making, were also classified according to preferences on hypothetical gambles as risk-averse or risk-preferenced. Scenario events were generated by a military expert, and decisions were evaluated with respect to the "ground truth" which he dictated. Subjects generated likelihood ratios representing the estimated impact of indicators on situational conditions rather than using externally provided likelihood ratios.

The results of this evaluation indicated an interaction between risk attitude of subjects and usefulness of decision aids. The performance of risk-averse individuals was found to be improved by use of the DDI aid, while performance of risk-preferenced individuals was not influenced significantly by use of either aid. No strong support was obtained for the idea that risk information like that offered by the NPRDC variation of TACAID was beneficial, although the possibility that alternative presentations of risk information or alternative task requirements might produce such benefits was recognized.

Two significant problems with the Bayesian updating component of TACAID were discovered, however, which constitute important qualifications on all other results. These problems relate to assumptions required by the aid that indicators are probabilistically independent of one another and that the situation being assessed is not changing in time (i.e., is stationary). In fact, tactical situation indicators typically exhibit conditional dependencies in that many

indicators may be quite redundant with one another so that observation of one will reduce the significance of the others. In order for Bayesian updating to be useful for such situation assessment applications, it will be necessary to develop a mechanism to effectively obtain and process conditional dependencies between indicators. The problem with the stationarity assumption is that the tactical situation may, in fact, change at any time, and prompt recognition of such changes is critically important. In several problems presented in the experiment, it was clear to the subject that the situation had changed long before the aid indicated the occurrence of the change. When indicators have led to the determination of a high probability for one condition, a great deal of new indicator information is required to revise that probability in favor of another condition according to a standard Bayesian updating process. A solution to this problem may lie in the development of a technique to use indicators separately to assess the possibility that the situation has just changed and then subsequently to assess the state of the current situation.

3.9 SRI INTERNATIONAL

SRI International accomplished four district project efforts for the ODA program. First, they investigated the decision environment of the Naval task force commander and identified appropriate types of decision aids for that environment. Second, they generated realistic, unclassified scenarios to be used in the conception, development, and evaluation of all decision aids in the ODA program. A third effort addressed the development of an aid for strike planning called the Strike Outcome Calculator (SOC). A fourth effort concerned the development of a general decision analytic aid based on the concept of the decision tree for the structured analysis of decision problems.

In investigating the task force decision environment, SRI examined how a variety of informational, organizational, physical, and personal factors influence the decision tasks encountered by the task force commander and his staff (Payne et al., 1974). Based on the characterization of decisions into the categories of planning, execution, and training, it was recommended that aids for planning decisions would hold the most promise for the ODA program, particularly in providing suitable applications for decision analytic techniques.

A collection of 36 generic decision-aid concepts was generated as candidate approaches to the difficulties identified in the decision environment (Payne et al., 1975).

Two scenarios were produced by SRI to provide a realistic context for developing, demonstrating, and evaluating decision aids. The two scenarios are the ONRODA Warfare Scenario (Payne and Rowney, 1975) and the Amphibious Warfare Scenario (Rowney, 1975). Both scenarios are embedded in the same hypothetical geopolitical situation, with the Amphibious Warfare Scenario occurring some years after the ONRODA Warfare Scenario. Most of the ODA decision-aid development projects used the ONRODA Warfare Scenario as the primary context for research efforts. That scenario postulated two Blue task forces with a mission to protect Grey from air assaults launched by Orange from ONRODA Island 300 miles from the Grey coast. Air strike warfare figured prominently in this scenario, with Orange launching strikes against both Grey and Blue task forces and Blue launching strikes against the Orange air base on ONRODA Island.

The Strike Outcome Calculator (SOC) was developed by SRI (Rowney et al., 1977; Garnero et al., 1978a; Garnero et al., 1978b) as an example of the kind of special-purpose outcome calculator which could be warranted for well-structured problems which are encountered fairly frequently. The SOC enables the analyst to describe own and enemy force plans in a strike warfare scenario and then to observe simulated outcomes. The aid requires the analyst to provide scenario and planning inputs, via an interactive computer terminal, to a fixed sequence of data tables. The required data characterize platform numbers and capabilities, repair and replenishment capabilities, weather forecasts, planning contingencies, etc. The aid is designed particularly for multi-strike planning applications with a fairly granular representation of individual strikes (e.g., a simulation time step of three hours) but with the capability to represent events and plans by which a later strike may be influenced by the results of an earlier strike. After the user has provided all necessary inputs to the SOC, output tables are presented which describe the simulation results for the defined strike plans. It is intended that the aid would be used iteratively to refine a plan until desired or acceptable results are achieved.

The Decision Structuring Aid is a general-purpose decision aid based on the concepts of decision analysis (Merkhofer et al., 1975; Merkhofer et al., 1977; Merkhofer et al., 1979; Merkhofer and Leaf, 1981). It guides the user through the construction and analysis of a decision tree which represents an unfolding decision situation. Each node in the tree represents a choice point for the decision maker or some external agent (e.g., enemy, nature), and the branches emanating from the node represent the feasible options at that choice point. Values are assigned to the terminal nodes of the tree (i.e., to the end results of the decision-event sequences), and probabilities are estimated for the branches emanating from the external agent nodes. Given these inputs, the aid "folds back" the tree to determine the expected value of each decision option available to the decision maker, thus providing a complete optimal plan for resolving all represented decision problems. Prompting and sensitivity analysis features are incorporated in the aid to facilitate complete, precise problem specification. It is intended that the Decision Structuring Aid would be used in conjunction with a variety of auxiliary aids which would manage data, calculate outcomes, and aid in the encoding of values and probabilities.

3.10 THE WHARTON SCHOOL OF THE UNIVERSITY OF PENNSYLVANIA

The Decision Sciences Department of the Wharton School of the University of Pennsylvania provided an experimental test bed for the ODA program (Oppenheim et al., 1979) and undertook the development of several information management aids. Although none of these efforts focussed specifically on the problems of strike planning, all of the information management aids are potentially applicable to a comprehensive strike planning system.

One aid is the Decision Aiding Information System (DAISY), which provides a user-friendly interface to a complex, structured data base (Hurst et al., 1975 a, b, and c). Another aid is the alerter, which is a program in a data base management system (like DAISY) which can be set by the user to generate alerts whenever particular conditions of special interest to the user obtain as the data base is dynamically updated (Buneman and Morgan, 1977; Ribeiro, 1978). An aid to the management and coordinated implementation of decision-aiding models was also conceived by the Wharton investigators

(Mitchell, 1976; Elam, 1979). Since it is clear that any comprehensive strike planning system must incorporate many models and a large dynamic data base, aiding techniques such as those developed at the Wharton School are proper components of the system.

3.11 OTHER ODA PROJECTS

For the sake of completeness, it is appropriate to mention three other ODA projects that did not concern techniques for design and evaluation of decision aids relevant to strike planning but did provide important guidance to the ODA program. CACI investigated the potential impact of decision-aiding systems on the organizational structure of decision-making functions in the task force (Spector et al., 1976). System Planning Corporation prepared a review of decision-aiding projects that preceded the ODA program (Lucas and Ruff, 1977). Finally, H.W. Sinaiko (1977) conducted a thorough interim review of the ODA program.

4. ISSUES IN STRIKE PLANNING AID DESIGN

4.1 PROBLEM IDENTIFICATION

The first step in decision-aid design is the identification of a problem in the decision-making process. It is generally difficult to define the problem to be addressed by a decision aid in adequate detail to guide all design efforts because most decision aids are intended to be at least somewhat versatile and because they are intended to apply to problems of the future. Problem identification should always begin with a characterization of the current decision-making situation and the problems that are currently evident. But it should be recognized that fleet inputs regarding the character of current decision problems will probably be derived from experiences with recent warfare or exercise situations and may not reflect adequate consideration of other current and future possibilities. Thus current problems should be extrapolated into the future to represent a best estimate of the problem domain at the time when the aid could be introduced for operational use. It is also appropriate to develop some specification of the diversity of problems to which the aid should be applicable so that all intended applications are considered in the design effort. Representative scenarios can then be constructed to provide a context for development and evaluation efforts. While these scenarios will then represent more or less contrived problems, it is also important that they still address the kinds of issues that are of current concern to the operational community.

4.1.1 Problem Types

The existence of a problem means the existence of some state which is not ideal. It is the objective of the decision aid to move the decision-making process from its current state to a state closer to the ideal along one or more

dimensions. The two basic dimensions for this decision-making space are workload and decision quality.

4.1.1.1 Workload. Workload is the product of time and individuals required in the decision-making process. A decision aid may benefit the process by reducing the time involved in the process, by reducing the number of individuals needed, or by simultaneously reducing both factors. Air strike planning workload is very high and allows much room for improvement. Large numbers of individuals are involved in the preparation of information for use in the process and in the decisions of the Strike Planning Board and the Strike Planning Team. The Strike Leader may spend hours (or even days) in the development of the plan.

Another factor in evaluating workload is the importance or skill level of the individuals involved. The time of a Task Force Commander is much more valuable than the time of another officer. Therefore, a more sophisticated interpretation of workload weights the time (or effort) of higher-level officers more than the time of lower-level individuals. Thus a decision aid may provide a greater benefit by reducing the required involvement of a Task Force Commander than by reducing the required involvement of other officers. The ODA program plan recognized this beneficial increase and initiated its efforts at the Task Force Commander levels.

In a similar vein, benefits with respect to workload may be achieved by constructing a decision aid which allows a lower-level individual to perform the task of a higher-level officer with equal competency and in equal or less time. Such an aid would allow the lower-level individual to replace the higher-level officer for that task in the decision-making process. The higher-level officer is thus allowed to concentrate on other (probably more critical) tasks.

A secondary benefit may be derived from such an aid in the reduced amount of training and individual capability required to perform the task. The aid may provide enough support and be easy enough to use that the training required to perform the task may be reduced. In addition, the capabilities of

the aided individual performing the task may not be required to be so great as they would be for unaided individuals. Thus a decision aid may result in a greater number of individuals being capable of performing a task and in these individuals being able to assume their responsibilities more quickly.

4.1.1.2 Decision Quality. Decision quality is a measure of how "good" a decision reached through the decision-making process is. A decision aid may benefit the process by increasing the quality of decisions made in the process. Unfortunately, while workload may be measured objectively through timing measurements and counts of personnel, decision quality is a strongly subjective measure. Decision quality is evaluated by considering (1) the spectrum of possible results of the decision aid, (2) the probability that each result occurs given that the decision is made, and (3) the value of each result. (This analysis of decision quality was investigated by Merkhofer and Leaf (1981).) Rarely may these aspects of decision quality be objectively determined. The probability of each result may only be estimated in most cases and may be subject to vast disagreement among experienced personnel. Similarly, the value of each result often produces different opinions from different experts. For example, in air strike planning the value of a result of two targets destroyed at a cost of three planes is definitely higher than the same result with the loss of an additional plane; however, disagreement may result in comparing the values of a result of one target destroyed at a cost of three planes and the result of two targets destroyed at a cost of seven planes. In addition, once the values and probabilities of the individual possible results of the decision have been considered, these factors must be used to compare this decision with another decision. The rules used for such comparison are also subject to much disagreement. As can be seen, evaluation of decision quality can be a difficult and controversial process.

So far, all that has been considered is the evaluation of decision quality for a given decision in a given situation. To evaluate the decision-quality benefits of a decision aid, the quality of decisions expected through use of the aid must be compared with the quality of decisions in the unaided

decision-making process under a variety of situations or scenarios representing the spectrum of situations encountered. Again, probability of each situation occurring may also be considered.

To establish its potential worth, the decision-aid design must be able to increase decision quality in at least one of the situations. This increase in decision quality must be achieved with no increase in workload or with an increase which is acceptable for the gain in decision quality and is also feasible within time constraints. So workload and decision quality must always be considered in tandem, not separately.

Because of the subjective nature of decision quality, it is often difficult to persuade a sponsor of the benefits of an aid addressing decision quality. Typically, a decision-making process is seen as being acceptable in its nature, with improvement needed only to make it faster. Greater imagination is necessary to envision the possibility of optimizing the decision itself as well as the time taken to reach it.

4.1.2 Problem Time Frame

In addressing the problems to be solved in the decision-making process, a time frame for the process must be set. Typically, the time frame is the present. In this case, an examination of the decision-making process as currently performed is made. The process is observed to identify areas of high workload and/or of poor performance or decision quality. Decision-aid designs may then be produced to improve the process in those areas.

However, the dynamic nature of the technology of modern warfare often makes such evaluations outdated not long after the studies producing them have been completed. Thus by the time the decision aid has been designed, implemented, and tested, the nature of the problem it is addressing may have been exacerbated to a degree that the decision aid may not prove significantly beneficial or, at the other extreme, the problem may no longer exist. So before proceeding with a decision-aid design to handle current problems, a study

should be made of changes envisioned in the problem between the present time and the expected introduction of the aid.

It is also possible that no problem is found in the current process, but that such a study would uncover problems to be encountered in the future. A current example of such a phenomenon arises from the prediction of the reduced availability of highly-trained decision makers among future Navy personnel. Thus where no need for decision aids is currently foreseen, decision aids may be required so that future personnel may perform at least as well as current personnel.

4.1.3 Decision-Aid Users

Identifying the potential users of a decision aid is done in one of two ways. The standard procedure is to identify the problem areas of the decision-making process with respect to decision quality and workload. Then the potential users are the personnel involved in that area of the process. For example, if routing an air strike is found to be a problem area, the natural potential users for an air strike routing aid are the Strike Leader and the Navigation Officer.

Another approach is used if the goal is to relieve an individual or group of its responsibility in the process. The potential user is then an individual or group capable of interacting with a perceived decision-aid design at a performance level equal to or greater than the original decision maker. As an example, the Strike Leader may be overburdened with responsibility in air strike planning. To relieve his workload, the responsibility for determining the feasibility of weapon loads on the plans to be used may be assigned to another individual with less training in the problem area, which may be compensated by a decision aid to help solve the problem.

To optimize the interaction between man and machine, attention must be given to characteristics and capabilities of the potential user. The command level of the intended user is especially important in determining the level of

detail to be addressed by the decision aid. In air strike planning, both the Task Force Commander and the Strike Leader may need an aid which predicts expected strike outcomes; however, since the Task Force Commander takes a broader overview, the aid for the Task Force Commander need not be as detailed as that for the Strike Leader.

Command level is also highly correlated with the importance of decisions being made. In air strike planning, the major decisions are made at the Task Force Commander level; the remainder of the planning may be complex and difficult, but it is no more than the detailed explication and optimization of a plan already determined by decisions made by the Task Force Commander. Thus although providing decision aids at the task force level may be difficult, any benefit derived is magnified by its major influence over the entire planning process.

4.2 COMPONENT DECISION-MAKING PROCESSES TO AID

As described in Section 4.1, when designing a decision aid a problem area in the decision-making process must be identified and the goals of the decision-aid introduction must be explicitly defined. Next, a more detailed consideration of the aspects of the problem area to be addressed must be made. For example, the problem area identified in the strike planning process may be routing the strike, with the goal of the proposed aid being improvement in the quality of routing decisions. Given this general outline, the decision-aid designers must then investigate the aspects of route planning where aids may increase the quality of the decision. Such aspects may include assessing the situation in which the strike is to be initiated, generating options in the routing plan, predicting possible outcomes of using different route plans under different assumptions, and evaluating and comparing the possible outcomes and their chances of occurrence. The following subsections describe these aspects, their instances of occurrence in air strike planning, and the degree to which they were addressed in the ODA program. A concluding paragraph discusses cost-effectiveness of the aspects that were addressed in ODA.

4.2.1 Situation Assessment

The first step in any decision-making process is to assess the situation presented to the decision maker. Each situation is a point in a complex multi-dimensional space. The situation is characterized by:

- the goals of the process,
- the resources available to the decision maker to achieve the goal,
- the goals of opposing forces,
- the resources of opposing forces, and
- neutral factors under the control of no forces.

The goals of the process are needed by the decision maker to provide direction to the process. One of the basic goals in air strike planning is the destruction of enemy resources. However, since air strike planning consists of numerous subproblems, relevant goals may also include the development of a strike route to minimize the duration of the strike when the strike force is detectable or a route which allows the least number of enemy aircraft to be deployed against the strike. Goals may include constraints such as maximizing enemy destruction within the constraint of a maximum acceptable aircraft loss or achieving a goal within a time limit. Goals may be stated in terms of probability (e.g., develop a strike plan with a 90% chance of destroying a bridge) or in terms of expected value (e.g., develop a strike plan expected to destroy 25 aircraft on the ground). As can be seen from these examples, strike planning goals are typically easy to express explicitly in a quantitative fashion; this feature is helpful in designing computerized decision aids.

The resources available to the decision maker are the instruments to be used in achieving the goal. The decision maker must know the types of resources available, their characteristics and capabilities, the quantity and location of each type, and their time schedule of availability. Depending on the level of the problem addressed, resource types may range in detail from attack

squadrons to specific attack planes loaded with specified weapons, ECM gear, and extra fuel tanks. Resources also include ships, fighters, AEW aircraft, SAR aircraft, ECM aircraft, and their fuel and weapon supplies. The great variety of resource types and relevant-type characteristics presents an overwhelming amount of data to be presented to the decision maker. This problem is reduced by examining strike planning in different levels of detail. A computerized decision aid may be able to handle such large amounts of data more quickly and with greater consistency than a human decision maker.

Understanding the goals of forces opposing the decision maker is key in the decision-making process. Often these goals are not known explicitly, although reasonable inferences may be reached about them. In strike planning, the goals of the opposing forces are typically clear (i.e., detect and destroy the strike force as soon as possible, conceal and evacuate as many valuable targets as possible, and destroy as much of the strike force as possible after the strike). As a result, such goals may be represented as relatively standard rules in computer decision aids.

The resources of the opposing forces must be evaluated in two senses: first as targets in the establishment of goals, and second as tools of the opposing forces intent on thwarting the plans of the decision maker. In strike planning, the opposing forces have resources similar in nature (if not in capability and numbers) to those of the decision maker, such as fighters, ECM aircraft, AEW aircraft, and ships. Radar sites, missile sites, and anti-aircraft defenses are resources valuable to forces opposing a strike, but not to strike forces. Bridges, troops, airfields, roads, fuel depots, and ports are resources important as possible targets, although not as tools to thwart a strike. The decision maker must know as much as possible about the opposing forces' resource types, capabilities, quantities, and locations. However, much of this information as gathered by intelligence is imprecise and often dated. The decision maker, before using the information, must first evaluate its worth, reliability, and timeliness. The analysis of such intelligence information is by itself an important area of investigation of computerized decision aids.

Neutral factors under the control of neither the decision maker nor the opposing forces must also be taken into consideration. In strike planning, weather conditions, terrain, and political boundaries are neutral factors which may have considerable bearing on the strike planned. Weather is an example of a random neutral factor. As in intelligence information, the variability of such a factor makes it difficult for assessment by a decision maker in many instances. Operational constraints imposed upon the decision maker by higher authorities may also be included in neutral factors. In strike planning, the decision maker may be constrained by standard EMCON and SAR policies.

Situation assessment is both time-consuming and difficult in strike planning. Great effort is devoted to collecting information from its various sources (e.g., personnel, manuals, data bases, reports, computer listings). The amount and diversity of information to be digested is overwhelming. Improvements in collection time and in guidance to information evaluation are needed.

Several ODA products offer components that are relevant to situation assessment problems, although none of the efforts addressed the high-level problem of collecting, integrating, and filtering information to construct a concise interpretation of a complex situation. The DDI TACAID demonstrated a feasible application of Bayesian updating for processing of dynamic indicator data to assess a tactical situation, although technical innovations to deal with conditional dependencies between indicators and non-stationary situations are required. The DAISY interface to a tactical data base and alerting techniques developed by the Wharton School are information management aids that can be helpful in assessment and monitoring of complex, dynamic situations.

Graphic aids for the presentation of summary information concerning a tactical situation were developed in many of the ODA efforts. DDI offered the triangle display of state probabilities relative to discrete decision thresholds. NPRDC conceived an alternative to the DDI technique for presentation of risk information in addition to state probabilities and decision

thresholds. ISC developed color-coded contour diagrams for displaying probability of aircraft detection by a field of enemy sensors. Analytics devised color-coded, bar graph displays for statistical distributions characterizing force readiness estimates. Grumman described a variety of map-based displays of tactical information, including summaries of combined threats and of force projection capabilities. DSA also offered a variety of map-based displays which convey information about combined sensor capabilities of a task force and vulnerability of the task force to enemy attacks.

4.2.2 Generation of Options and Hypotheses

Once an assessment of the situation has been made, the decision maker must generate hypotheses and conclusions about the interactions of the resources available and use these conclusions to generate a set of options. This aspect of decision making requires creativity and imagination in conjunction with experience and historical perspective. In strike planning, options considered depend on the problem and the level of detail required.

In strike planning, options are the combinatorial instantiation of:

- quantities of aircraft and weapons;
- routes, schedules, stations, and rendezvous points;
- policies and practices such as EMCON;
- targeting details; and
- alternate plans.

The total number of possible options is combinatorially an overwhelming number in strike planning. The total possible combinations of aircraft, weapons, routes, and schedules is too much for any decision maker, man or machine, to consider. Thus decision makers must concentrate on boundary or limiting situations (e.g., using all aircraft available, using most powerful weapons available, using most direct route) and on representative options (e.g., single

numbers to represent small, medium, and large quantities). Even at this reduced level of detail, the number of options may prove too large to preclude confusion on the part of the decision maker. Heuristics may be used to identify only those options with a reasonable chance of being acceptable or being improvements over other options; however, heuristics may overlook some surprisingly effective options unless the heuristics are very sophisticated (which then often negates their primary strengths -- speed and ease of computation).

The hypotheses in strike planning generally concern the actions taken by the enemy upon discovering the impending strike. Thus hypotheses are generally the generation and evaluation of options open to the enemy to better evaluate the options in the strike plan. These hypotheses are not as complete as the strike options. However, they often involve some added probabilistic evaluation of occurrence likelihood.

Several ODA projects developed innovative techniques for option generation. The ISC work in route planning offered display aids for manual generation of routing options and optimization algorithms for automatic determination of locally optimal routes. The Decision Structuring Aid developed by SRI provided guidance in the generation of decision options through structuring of problem representation, prompting for consideration of significant conditions, and sensitivity analysis to determine decision tree branches where options should be further developed and analyzed. The displays and analyses offered in DSA's EWAR offered a variety of types of assistance in the generation of task force EMCON plans (e.g., by displaying plausible enemy inferences as to identities of ships).

4.2.3 Prediction of Outcomes

Given the strike options and the hypotheses about enemy response, the decision maker then matches the options with the hypothesis to produce predictions of the outcome. For each option-hypothesis combination, a set of outcomes may be predicted, each with a probability of occurrence. Using the probability of each hypothesis actually occurring, a new aggregate set of probabilities may

be attached to each outcome in the set for a given option. Each outcome in strike planning could include:

- quantities of each type of aircraft lost on each side,
- quantities of each type of weapon expended,
- targets destroyed or incapacitated, and
- objectives gained in a ground action supported by an air strike.

This aspect of decision making received much attention in the ODA program. SRI produced a Strike Outcome Calculator (SOC) which, given the participants in a strike (or campaign of strikes), calculated the result of the strikes. Analytics also predicted strike outcomes in its Air Strike Timing Decision Aid (ASTDA). These prediction models were not constructed with the ideal of actual prediction, but were rather given only a semblance of reality to produce results to test more important features of the aids. The argument given was that the place of the models used could later be filled by more accurate predictions of outcome. Unfortunately, when these aids were evaluated by Navy personnel, interest focussed on the prediction models rather than the other features which the models were intended to highlight. Thus although much effort went into research in prediction models, not enough was done to validate them or even to find how to validate them. Much of the realism of these models depended on the validity of the data input to them; it was also questionable whether "valid" inputs needed by the models could be reliably obtained in combat situations.

4.2.4 Evaluation of Outcomes

The comparison of outcomes (whether predicted or actual) is not as simple and direct as it would seem at first glance. Typically, positive results in an outcome are balanced to some degree by negative results in other aspects. The counterbalancing trends in the results can be confusing to a decision maker. Often the decision maker must focus attention only on selected aspects with large effects and disregard the rest to reach a conclusion. The difficulty is

compounded when the decision maker considers that the outcome is only one of many possible outcomes from a selected option, and to correctly evaluate the outcome of the option, some probabilistic distribution for the outcomes must be considered in the overall evaluation.

Strike planning has many problems in the area of outcome evaluation. Each outcome is multi-dimensional in nature, with the dimensions not being easily comparable. How does a decision maker value the air support in a successful land operation or the interdiction of supply and reinforcements? How many aircraft would the decision maker sacrifice to destroy a target? The answer may be quite different for various types of aircraft, targets, and tactical situations. In such multi-dimensional outcome situations, not only is it difficult to balance the value of the dimensions, but it is also hard for the decision maker to remain consistent in the valuation of the dimensions throughout the process of outcome evaluation. Decision aiding in outcome evaluation can provide a consistent system and quantitative way of providing the decision maker with a baseline evaluation of strike results.

Value models were incorporated in most of the ODA-sponsored aids. Models for valuation of states or events pertinent to the tactical situation were employed in DDI's TACAID, Grumman's Options Selection Checklist, and SRI's Decision Structuring Aid. All of these aids used linear, additive forms for calculation of multi-attribute utilities. Value models (also linear and additive) based on predicted outcomes rather than generic states were incorporated in Analytics' ASTDA and DSA's EWAR. The route planning aids investigated by ISC employed a non-linear value model based on variables which are assumed to exert significant influence on the situational outcome (i.e., fuel utilization and probability of detection by enemy). Although evaluations of the ODA products disclosed some resistance by experienced Naval decision makers to use of value models in operational aids, it is likely that properly designed and implemented value models will ultimately be accepted. Training programs and supporting aids will be required to help in elicitation of component utility assignments and in verification of the consistency of those assignments.

4.2.5 Implementation of Selected Action

Even after a decision has been made, the work of the decision maker is not completed. The decision maker must place the decision or plan into a form suitable for its implementation into action. In the case of strike planning, numerous schedules, requisitions, and orders must be submitted to put the strike plan into action. This work, though not difficult, is time-consuming. One of the indirect benefits of constructing a decision aid for strike planning is that the data needed to issue the required paperwork is stored in the computer as a consequence of using the machine in the decision process; the required paperwork may then easily be composed by the computer decision aid with relatively little additional programming work.

The ODA program put little effort in this area. Since the capability of the computer to perform this function is already recognized, no research was actually needed. Still the SRI Strike Outcome Calculator demonstrated some capability in this area with its displays of aircraft assignments and strike schedules for a campaign. Minor modifications to EWAR (DSA) and ASPS (ISC) could easily be envisioned to produce the paperwork for implementing the decision selected. However, development of such details of aid operation are probably best postponed until the aid has been accepted for introduction into the fleet.

4.2.6 Processing Feedback from Selected Action

Decision makers must take advantage of any information, results, and conclusions from actions taken as a result of their decisions. This recent data must be integrated with the data collected in the past so that future decisions can be made with the most recent information. The information may also provide alterations to the decision maker's internal models of the capabilities and interactions of resources. In strike planning, such information may include:

- the outcome of the strike,
- reports on the capabilities of aircraft (friendly and enemy),

- reports on enemy resource capabilities (e.g., missiles, anti-aircraft capability),
- reports on numbers of enemy forces encountered,
- reports on enemy responses to attack,
- reports on enemy tactics, and
- reports on enemy detection capability (e.g., how soon was the strike detected?).

Much of this information could be integrated into a data base. Some of the information could only be implemented by modifying the models used; this aspect is more difficult.

The ODA program did not pursue the issues in this area. This policy was reasonable since the program was performing research in how the initial decision was to be made. The feedback issues can be better addressed once the initial decision is well-understood.

4.2.7 Cost-Effectiveness

When designing a decision aid, consideration must be given to the cost of developing the aid, the cost of using it, and the cost of maintaining it. Decision-aid designers too often tend to stress the benefits and improvements to the decision process without also considering their attendant costs. Table 4-1 examines the relative benefits and costs to be derived from decision-aid development in different areas for strike planning. Situation Assessment and Evaluation of Outcomes have high benefits since their concepts are clearly understood and the basic improvements in decision quality and speed are significant. Generation of Options and Hypotheses, Processing Feedback, and Prediction of Outcomes are all areas of great value; however, the current level of conceptual knowledge in these areas is not as high as in others. Thus the benefit at the present stage of knowledge is only moderate (i.e., great potential with only limited degree of success). Implementation of Selected Action

Table 4-1. Benefits and Costs of Decision-Aid Areas in Air Strike Planning

		COST		
DECISION-AIDING AREA	BENEFITS	DEVELOPMENT	EXECUTION	MAINTENANCE
Situation Assessment	High	Medium	Medium	High
Generation of Options and Hypotheses	Medium	High	Medium	High
Prediction of Outcomes	Medium	High	Medium	High
Evaluation of Outcomes	High	Medium	Medium	Medium
Implementation of Selected Action	Medium	Low	Low	Low
Processing Feedback	Medium	Medium	High	High

is an area well-conceptualized, but the benefit to be derived is only one of reduction in work; no benefit to decision quality is perceived.

In analyzing cost, Generation of Options and Hypotheses and Prediction of Outcomes are areas requiring much research and precision for acceptability; hence the high development costs. However, the cost of execution in these areas is medium as a result of high computer processing costs combined with a low cost in human participation. Maintenance in these areas would be a high cost to adapt the models to the ever-changing world of air warfare. Situation Assessment and Evaluation of Outcomes are of similar nature in cost. In each, the development need not be of an advanced ground-breaking nature but would still be time-consuming. The aids in these areas are true man-machine interactive aids; the aids do not replace, but rather enhance, the decision maker's capability. Hence the execution of these aids entails the cost of a decision maker's extensive participation, but with relatively low processing cost. Maintenance for Situation Assessment may be high to acquire and store needed information in the data base. Maintenance of Evaluation of Outcomes is less since the amount of data required is not high; still the models used for evaluation will require changes to suit the situation and the opinions of the decision makers. Processing Feedback requires high cost of maintenance just as Situation Assessment does since it is a form of Situation Assessment (i.e., updating the situation). Since its execution is maintenance, execution is also high. Techniques to perform Processing Feedback are understood well, but their optimization would make development a medium-cost task. Finally, Implementation of Selected Action is of low cost in all areas. Development of the software to issue reports and orders is well-understood and relatively inexpensive. Execution is automatic once the plan is set. Maintenance of data necessary for the job is actually handled in Situation Assessment. Thus Implementation of Selected Action is of low cost once other aiding areas have been developed.

Table 4-1 supports the ODA program's emphasis on Evaluation of Outcomes. The benefits to be gained from this research outweighs the expected cost. Situation Assessment and Implementation of Selected Action appear to be areas of promise for future development.

4.3 ALLOCATION OF FUNCTIONS BETWEEN MAN AND MACHINE

Once the designer has analyzed the cost and benefits of aiding various aspects of the problem area, a more detailed examination of the degree of aiding required must be made. Generally, the approach made is to allocate as many functions as possible to the aiding device in order to reduce the workload of the decision maker or even free the decision maker for another task. There are, however, several significant considerations that constrain the level of participation of the aid in the process.

The most important of these considerations is the quality of aiding that the machine may provide. For many functions, a man is equal or superior to a machine in the quality of the decision made. The best examples of this case arise in decision situations which require creativity or insight. In air strike planning, an example is a situation where the decision maker must outline the tactics for the attack on the targets (e.g., the number of waves and their synchronization). A machine aid currently may be capable of an acceptable standard solution, while an experienced human decision maker is capable of valuable insights which may significantly improve the quality of the decision.

Another constraint is the time needed by the machine to reach a decision which is better than or equal to that of a man. Such situations often arise where optimization is required. The machine may have to consider an exhaustive number of cases to reach the best conclusion, while a man may be capable of quickly reaching a conclusion which is virtually as good as that reached by the machine. A possible example of such a situation in strike planning is weapon loading, where a machine may have to go through many combinations of weapons, aircraft, delivery option, and targets to reach an optimal conclusion, while a human decision maker may be able to reach a near-optimal conclusion in comparable or shorter time.

It has been thought that the decision-aid designer must allocate the processes to be handled by the human and the machine according to what each does best. Several questions arise here. First, what is meant by "best" depends on the goal of the decision aid. If the aid is meant primarily to reduce the

workload of the decision maker, a function should be assigned to the machine if the machine can perform the function in time and with quality comparable to that of the human decision maker. If the goal is to improve the quality of the function's result, the function should be assigned to the machine if it significantly improves decision quality without incurring an unacceptable penalty in time taken to perform the task. Similarly, if the goal is to decrease the time taken to perform the function, the function should be assigned to the machine if the assignment decreases time to perform the function without incurring an unacceptable decrease in decision quality.

Second, this allocation process cannot be performed simply by considering each function independently of the others. If each function is assigned in this manner, the decision-aid designer may seriously underestimate the time required for the entire process to be completed. The reason is because this independent allocation process does not account for the time and processing effort needed for interaction between man and machine. If two sequential functions are assigned to one system, man or machine, little if any data transfer is required between the two processes since most of the data needed for the second function would already reside in the system as a result of processing the first function. If the two sequential functions are assigned one each to man and machine as a result of analyzing which system performs the function best, it is likely that the system performing the first function may have to communicate a great deal of information to the other system so that it may perform the second function. Such communication between man and machine can be cumbersome and time-consuming enough to negate the advantages gained by the allocation of the best system for each function. Thus keeping interaction between systems to a minimum is another consideration.

The decision-aid designer, besides allocating the functions to create the decision, must also allocate the responsibility for their implementation. If a machine aid is selecting a decision option, the entire decision-making process is speeded and made less costly by allowing the aid to directly implement the decision by issuing the orders to the units participating in the action.

However, this procedure allows for no human participation in the decision. Only in circumstances where time and manpower are extremely limited is it likely that the human decision maker would allow such a procedure to occur. In most situations, the human decision maker would like at least to review the decision formulated by the aid and would also probably like to examine the process by which the decision was attained. For strike planning, this is analogous to the review of the final strike plan by the Strike Planning Board.

The problem with such review and interaction is that it slows the process and requires extensive participation of a human decision maker, depending on the level of detail required by the decision maker and the confidence the decision maker has in the aid. If the required level of detail is high and the confidence is low, the monitoring of the aid can negate the aid's value by reducing its cost-effectiveness below the point where its benefits outweigh its cost. Such a case can occur where the aid itself is monitoring a situation and the decision maker must simultaneously monitor the aid since the aid is not sophisticated enough to handle all cases with sufficient confidence. Fortunately, strike planning is not such a case. It appears possible to allocate substantial functions of the strike planning process to machine aids, with the aid results requiring review only after termination of the function. Since continual monitoring of aids does not appear necessary, the applicability of aids to strike planning is increased.

It is also possible for the aid to monitor the human decision maker. For example, a human decision maker may be required to submit his decision to the aid so that orders may be generated for decision implementation. The aid could then check the decision against general criteria established at a prior time and veto the issuance of the orders if the criteria were not met. It is apparent that such an aid feature works well if the aid incorporates the Implementation of Selected Action aspect of decision making in its procedures. In strike planning, such a feature could also provide a benefit in reduction of workload of high-level personnel in that it could replace or reduce in time the work by the Strike Planning Board to review the final strike plan.

Allocation of functions or responsibilities to man and machine need not be fixed. The allocation may depend on the situation being faced. For example, in situations where time is at a premium and decision quality need not be optimized, functions normally allocated to man and machine together may be allocated to the machine aid alone. Similarly, if the human decision maker is not pressed for time, the man may assume more functions than normal. In this way, machine aiding could be used only when required, and the human decision maker could remain in touch with the aspects of the process he typically does not handle. It would be the task of the human decision maker to make this dynamic allocation of responsibility and function according to the circumstances of the situation.

The only direct investigation in the ODA program into the allocation of functions between man and machine was the research conducted by ISC on the route planning process. That effort examined the performance consequences of assigning the jobs of generating initial candidate routes and iteratively refining those routes to various computer algorithms or to a human decision maker. For the particular routing problems that were posed which were simplified versions of realistic strike routing problems, it was determined that initial candidate routes should be provided by the human and that subsequent optimization could be performed about equally well and equally rapidly by man or machine. It is important to recognize in this case, however, that the rate of current advances in micro-electronic technology and algorithmic efficiency, along with application to realistically complex route planning problems, will probably tip the scale in favor of allocation to the machine of at least the optimization function.

In the process of designing and creating decision aids, the allocation process itself was performed by the contractors in the ODA program. Areas of research were selected according to the benefits expected to be derived by having strike planning functions performed by a machine aid. In several projects, comparisons were made of human and machine performance to determine if

1) the machine performance was significantly superior to human performance,

2) machine performance was adequate to replace the human decision maker, and/or

3) machine participation enhanced the performance of the human decision maker.

Using the results of these efforts, an attempt was made in the ODA program to design a single aid integrating the ideas found to be successful (Glenn and Bennett, 1980). This integration process was another form of function allocation. Various machine aids were allocated to disparate problems in the strike planning process according to their capabilities relative to those of a human decision maker. The design was made flexible to allow dynamic allocation of roles by the human decision maker according to the circumstances of the situation.

Thus the ODA program did explore how aid performance could be evaluated against human performance to provide an experimental basis for allocating functions to man and machine. In addition, the program applied these results in an allocation of functions for a proposed Air Strike Planner (ASP).

4.4 OPERATIONAL CONSTRAINTS ON THE DECISION PROCESS AND AID DESIGN

Every decision process is composed of a number of general aspects and stages, as described in Subsection 4.2. These common aspects provide the framework which describes similarities among all decisions and how they can be aided by a common decision-aiding technology. However, it is also the case that no decision-making process occurs in a vacuum -- each decision process is situated in a tactical and operational context which includes a hardware system, a software system, an organizational structure, and a history, as well as a tactical situation. These factors act to constrain the decision-making process and to make each instance of a given decision different from each other one. These unique aspects or individual differences in decision-making processes thus provide a very crucial kind of input to the decision-aid design process, particularly for aids intended for operational use.

4.4.1 Hardware Constraints

The operational Navy has considerable constraints on its computational capability. These arise from the need to use only that equipment which can fit on-board, which has been adequately ruggedized, and which is compatible with other on-board equipment. For the most part, these computational constraints restrict the speed and complexity of the computations which can be carried out by decision-aiding algorithms. Thus computation-intensive approaches (such as the ASTDA outcome calculator or the EWAR algorithms) can quickly come into conflict with the limited processing capabilities of on-board operational platforms.

Most ODA projects did not limit themselves to current operational hardware constraints and tried to restrict space and execution time to realistic levels. This was appropriate given the program's orientation of technology development. It was also appropriate given the extremely rapid rate of advance in computer hardware technology since it is likely that algorithms which today require sophisticated, state-of-the-art mainframe computers will, in five years, be computable on stand-alone desk-top processors. Given these two concerns, it is reasonable that computational requirements have not played a major role in the choice of techniques for ODA project aids. Still it must be remembered that as decision aiding moves toward operational implementation, all aids must be designed to function within the computational constraints of the ultimate environment of use.

A second kind of hardware constraint is in the display technology area. A number of ODA projects (e.g., ASTDA, DAISY, EWAR) make use of color graphic display technology. While the fleet does have widespread availability of monochrome graphic displays, there is virtually no color capability, nor is any planned for the near future. It is therefore quite possible that many specialized interface techniques which rely on color graphics may prove inappropriate for fleet use except in the far term.

4.4.2 Software Constraints

Software constraints pose a similar though less severe problem for aid design. Many advanced software methodologies used in the ODA test bed, as well as many others used in ODA products developed outside the test bed, are not currently available in Naval operational computer systems. Some examples of these are:

- multi-job execution capability from a single peripheral device (central to DAISY),
- virtual memory (used by all test-bed-implemented aids),
- relational and network data base management systems (also used in DAISY), and
- APL interpreters (used in the Decision Structuring Aid and TACAID).

With the current effort to bring all DoD computing under the umbrella of the ADA super language, it is likely that these sophisticated capabilities will be slow to enter the operational environment and that many will not make it at all. Thus the dependence of decision aids on state-of-the-art software technology may also be a hindrance to operational usage.

The hardware/software environment combines with the operational setting to impose other more subtle constraints on decision-aid design and implementation. For example, data is available in an on-board computer system at only a given level of detail and in only a specific format. The portions of the fleet/platform organization which collect and process these data, as well as the software and hardware used for recording, collecting, and processing them, are fixed aspects of the environment, at least to the aid designer. Changes to these capabilities will probably be beyond the power of the decision-aid designer to implement. Thus the aiding algorithms will have to be designed to use the data that are "at hand" within the operational environment. A number of the aids developed in ODA have assumed that highly sophisticated kinds of data will be available (e.g., hour-by-hour probabilities of individual aircraft being ready for use), whereas in reality they are

not. When designing aids for operational use, such assumptions must be carefully avoided.

4.4.3 Organizational Structure and System History

System and organizational factors go far beyond the data availability level. The Strike Planning Board and the Strike Planning Team will have standard operational procedures for interfacing with the various other commands that supply them with information, coordinate resource utilization with them, implement their decisions, and so on. In addition, the team immediately involved in the decision has an internal structure which has arisen from standard procedures and the history of the organization and individuals involved. These organizational interconnections, both internal and external, necessarily exist in every military decision-making process and provide a far-reaching set of constraints.

The designer of a decision aid walks a fine line between designing a decision aid and designing a new decision process. In any system where a given decision is currently being made without a decision aid, there will exist a set of standard procedures and relationships which are utilized to make that decision. These procedures and relationships may or may not be optimal from the perspective of making an unaided decision, but they represent a framework of a decision process which the aid designer must consider. In particular, the aid must either be designed to fit within the existing organizational framework for the decision process or to require a totally new framework.

If the former approach is chosen, then the constraints on the design process are enormous -- every "player" in the decision process must maintain the same role and produce outputs in the current form from inputs in the current form. The prospects for significant impact on the decision-making system are diminished simply by virtue of the limited range of design possibilities. However, the prospects for an easy and efficient transition into operational use for such an aid are great. Little new training, outside that required for aid use, is necessitated, and the points of contact between the

decision-making system and the remainder of the organization are unaffected. Thus there will be little organizational resistance to introduction of an aid designed with this approach.

If, on the other hand, the aid designer chooses not to work within the framework of the existing organizational decision-making process, then an entirely different picture emerges. The designer in this approach is faced with the task of creating not just an aid but an entirely new decision process, one which includes a decision aid and a new set of relationships and procedures in the decision process. The possibilities of devising a system which produces significantly higher-quality decisions, which significantly reduces workload, or both, is greatly increased when the aid designer has such latitude. However, the scope of the innovation required in the existing organizational setting is also greatly increased with this approach. If a new set of relationships is required to utilize the decision aid or if a new set of procedures must be employed to make the decision, then substantial retraining of current personnel and substantial retooling of existing support equipment will be required. This will create two serious kinds of resistance to the introduction of the aid. One, of course, is cost -- both in terms of time and money. The more changes and retraining that are required to effect the aid's introduction, the greater the cost. The second factor is organizational. The modification of existing command and personnel relationships will create an instability in the overall system -- "make waves" -- and engender resistance from an organization which sees itself already making the decision with existing procedures and no obvious deficiencies. Promises of lowered workload and improved decision quality, unless they are convincingly substantiated, are unlikely to motivate organizations into wholesale changes in proven (if suboptimal) procedures.

Lest this discussion seem critical of the fleet, it must be pointed out that in organizations where instability, for even short periods of time, is potentially catastrophic (such as national defense organizations), the avoidance of such instability is understandable. The organizational impact of decision aids on organizational structure was investigated early in the ODA

program by CACI (see Spector et al., 1976). Their results indicate that the impact of aid introduction is highly dependent on the individual decision-making styles of the aid users and the history of the organization into which the aid is introduced. Thus it is difficult to predict the specific kind and extent of organizational changes that will be required by an aid in a given fleet environment.

4.4.4 Tactical Constraints

Strike planning is not a solitary activity that is conducted by the strike planning board and team in isolation, but is only one part of a complex network of decisions and responsibilities that are ongoing during fleet and task force operations. The same is the case for all other principal decision "nodes" in the command and control network. The interconnectedness of the overall C^2 process creates an interdependence among the individual decisions in the process, and this interdependence also constrains the individual decision processes (and decision aids for them). For example, in fleet operations the decisions made regarding antisurface warfare, antisubmarine warfare, and antiair warfare all impact the strike planning process in a variety of ways and vice versa, because all these warfare area commands compete for the same fleet resources.

Moreover, there is a temporal interconnection among decisions in the overall C^2 process. If a decision to launch a particular kind of strike at a certain time is made, then that constrains the kinds of operations which can be conducted before the strike force returns, defines the interval in which search-and-rescue must be anticipated, and so on.

The important lesson to be learned from this interdependence of command decisions is that it is always inappropriate to assume that the decision maker will have the ability to unilaterally request resources, gain access to needed information, or generally implement his plans as desired. Rather, the aid must be accommodated to the interactions with other decisions in the C^2 process which are contingent upon the decision addressed by the aid, and the aid must allow for the interplay of these decision processes. For example,

route planning is certainly a decision where decision quality can be improved by an aid, but the route planning process is one which impacts many other decisions and vice versa. The strike force composition, the search-and-rescue resources available, and intelligence data on enemy force readiness are all factors which may affect route planning decisions.

A more direct tactical constraint on the decision-making process is the time available. While most decision aids developed in ODA did not reflect a direct concern for time-constrained decision making, virtually every critical tactical decision is, in fact, time-constrained. To make matters more complicated, the time constraints vary from instance to instance of the decision. In strike planning, for example, the planning process will usually be undertaken 12 to 24 hours before the anticipated strike, with the final decision being required several hours before the launch time to allow for aircrew, aircraft, and carrier preparations. Thus several hours will usually be available for planning, more than enough time to use any (but perhaps not all) of the strike planning aids developed in ODA. However, there will be occasions in which a strike plan must be developed very rapidly for a launch time as soon as possible. In such instances, the use of aids would be especially beneficial, but if there is not an abbreviated manner in which the aid may be consulted, then it might have to be bypassed because of the time limitation. Virtually all of the ODA aids have a single, standard mode of use which requires a fairly fixed amount of time. While this induces a valuable consistency in the decision-making procedure, it is not responsive to the changing time constraints of tactical military decisions. For decision aids in operational use, the ability to use them in a variety of time-degraded modes is an important feature.

4.5 DECISION-AIDING TECHNIQUES

Given that there are a variety of parts of the decision process which an aid might address (e.g., situation assessment, outcome evaluation, information processing) and a number of problems which a decision aid might solve (e.g., overall workload, poor or inconsistent decision quality), it appears that there would be no one technique or method which can be appropriate for

all decision-aiding situations. Thus it is not surprising that there have been a large number of different techniques explored in the ODA program. Some of these have been new applications of existing techniques, such as mathematical programming, while others have been totally novel concepts, such as operator-aided optimization.

The ability of each technique to address different problems and parts of the decision process has not been systematically explored in ODA. Instead, applicability has only been demonstrated by example. Fortunately, however, the relationships among techniques, problem characteristics, and parts of the decision process can be assessed in a fairly direct manner. This subsection examines these relationships.

To facilitate this examination, the various decision-aiding techniques investigated during the ODA project are divided into four broad categories. The first is *interface/interaction* techniques and includes the various methods used to facilitate the transmission of information between human decision makers and computer aids. The second is *mathematical/analysis techniques* and includes methods for applying quantitative analysis to support the decision-making process. The third category can be called *"scientific guessing" techniques* and includes methods for inducing consistency across individuals and for structuring the decision-making process of human decision makers. The fourth category is *training/practice techniques* and includes methods for allowing human decision makers to develop expertise through experience in a given type of decision-making problem. The specific ODA use of techniques of each type is discussed below, and the general applicability of each type of technique is analyzed in the following paragraphs.

4.5.1 Interface/Interaction Techniques

The transfer of information from computer to human decision maker and vice versa is a crucial part of the decision-aiding process. To the extent that a decision aid is intended to relieve decision-maker workload in an information-saturated environment, it is essential that the aid make the appropriate information accessible to the human in as efficient a manner as

possible. Techniques for effecting an efficient flow of information between man and machine are called interface/interaction techniques.

A number of special techniques developed or advanced in ODA have focussed on this issue. The friendly interface to a complex data base developed as part of DAISY provides a way for a decision maker to quickly retrieve exactly the information he needs simply by asking the computer for it. The multiple-window terminal interface (i.e., "windowing"), also an adjunct to DAISY, provides a way for a decision maker to maintain communication with a number of important information sources at the same time from a single computer terminal. This technique shields the decision maker from data flow interrupts -- situations where he must divert his attention (and presumably his terminal) from one data source because another one is prepared to begin transmitting.

Another way that interface/interaction techniques can help relieve decision-maker workload and perhaps improve decision quality as well is to present information at a higher level than otherwise available. This type of information integration technique allows the decision maker to see the "big picture" more quickly and thus be in a better position to make a more informed decision. One ODA technique based on this approach is the triangle display technique used in DDI's TACAID, which displays a complex three-dimensional relationship in a simple manner. Another is the delta-biased uncertainty band display technique in Analytics' ASTDA. This technique presents the central tendency, variability, and skewness of a probability distribution to a decision maker in a single symbol so as to allow easy comparison of whole distributions. Still another technique is the trade-off table used in DSA's EWAR; this table presents the decision maker with the full range of trade-offs possible between two competing decision-making criteria.

A final way in which interface/interaction techniques may aid decision making is to improve decision quality by transforming a problem into a format which takes advantage of some specialized human information processing capability. The use of graphic displays to assist in the solution of spatial problems is an important area in which this approach is useful. The ISC route

planning aid maps and similar tactical situation displays in DSA's EWAR are examples of aid interface techniques which assist a decision maker in solving a problem by allowing him to use his visual processing capabilities.

In general, then, it appears that interface/interaction techniques are most appropriate for decreasing decision-making workload in situation assessment tasks and in improving decision quality in spatial analytic tasks associated with decision making.

4.5.2 Mathematical/Analysis Techniques

Many decision problems in strike planning involve manipulations of information which are conceptually simple (or at least tractable) but which are beyond the capability of human decision makers to carry out within the time available to make a decision. For such problems, mathematical/analysis techniques are useful. One of the most powerful techniques available within this class is mathematical programming. Decisions which require some form of optimization are often simple in *structure*, as in weapons allocation for example, with the difficulty arising only from the overwhelming number of alternatives which must be considered and/or the complexity of the calculations involved with evaluating and comparing the alternatives. Optimization techniques, especially mathematical programming, are useful for improving decision quality in these problems. Mathematical programming was investigated in ODA in the route planning research of ISC. In that research, aids were developed that optimized user-selected initial strike routes with non-linear programming and dynamic programming techniques.

Another set of techniques of this class which has great aiding value is outcome calculators. Four general kinds of outcome calculators can be built:

- Monte-Carlo simulation models,
- deterministic or mechanistic simulation models,
- probabilistic outcome calculators, and
- closed-form analytic equations (e.g., Lanchester equations).

Only the first two kinds were used in ODA; a Monte-Carlo calculator was used in ASTDA, and deterministic simulations were used in EWAR and SOC. The outcome calculator is primarily of use in improving the quality of decisions which involve prediction of future events. In cases where decision makers currently perform detailed manual calculations of the probable results of candidate plans, availability of outcome calculators could also serve to reduce workload as well as to improve decision quality. This appears to be the case in air strike planning, in which the planner must typically perform complex weaponeering, fuel utilization, and force attrition calculations based on data obtained from JMEM and NATOPS documents. In other cases where decision makers do not currently undertake detailed calculations of estimated outcomes but rather rely on experience and simple heuristics to make predictions, however, it should generally be expected that the interaction necessary to initialize and run an outcome calculator would tend to increase workload. Such increases may, of course, be justified in terms of the corresponding increases in decision quality.

The explicit modeling and treatment of uncertainty in decision problems is another way in which mathematical/analytic techniques can be beneficial. By quantifying the uncertainties and folding them into the expectations of results of actions (usually via outcome calculators), a broader and more accurate *range* of possible outcomes can be constructed and made available to the decision maker, thereby giving him a more complete picture of the possible future situations.

Mathematical/analysis techniques are thus generally useful for improving decision quality in problems where prediction is important or where complex combinatoric properties affect the decision process. They are also useful for reducing workload in situations where extensive manual calculations are currently required.

4.5.3 Scientific Guessing Techniques

It is now well-established that there are great individual differences among human decision makers and that human decision-making performance is

degraded in a variety of predictable ways under stress, fatique, and time constraints (see Slovic et al., 1979). It is also clear that there are certain aspects of decision making in which humans perform poorly, such as interpretation of very high or very low probabilities, compensation for correlation among indicators, or evaluation of additive linear utility functions. A great deal of suboptimal performance of human decision makers stems from the various limitations of human cognitive capabilities. A number of techniques have accordingly been developed to support these aspects of human decision making. These techniques can be called "scientific guessing" techniques because they are derived from mathematically-based models of normative decision-making, and their collective (and individual) goal is to introduce a kind of scientific rigor into the decision-making process.

One family of these techniques treats the processing of probabilistic information. Techniques such as Bayesian updating use theorems from probability theory to provide a way of systematically applying data on indicators, especially intelligence indicators, to modify probability estimates of uncertain events as new information becomes available. Bayesian updating is a key technique in the DDI TACAID.

Another type of scientific guessing technique is decision structuring. The technique of decision analysis was originally developed as a way of helping a human decision maker carefully and systematically structure the decision problem so as to avoid the omission of important aspects of decision options. Decision analysis is automated in the SRI Decision Structuring Aid and is a key part of the DDI TACAID. An alternative method for decision structuring has been developed outside of ODA by Pearl et al. (1980) based on the analysis of decision-maker goals.

Another group of scientific guessing techniques deals with the explicit representation of preferences. Multi-attribute utility theory has been developed to assist decision makers in making consistent trade-offs among competing criteria in complex decision problems. It was incorporated into the DDI TACAID, the SRI Decision Structuring Aid, the Grumman Options Selection

Checklist, and Analytics' ASTDA. Multi-attribute utility models generally assume that there is a well-defined linear trade-off possible between competing criteria at all points of the criterion measurement scale; in cases where this assumption does not hold, non-linear utility models may be more appropriate. Non-linear utility models generally incorporate non-quantitative heuristics of decision makers into a precise quantitative formula for evaluating alternatives quickly and consistently. The ISC route planning aids use a non-linear utility function to evaluate candidate ingress strike routes.

All of the various scientific guessing techniques address decision quality more than workload. Without such aids, experimental evidence suggests that humans make probability assessments and value judgments rapidly but inconsistently and poorly. Moreover, without structuring aids, decision makers tend to deal with an oversimplified internal model of the problem, but usually in a rapid manner. However, these kinds of aiding techniques have substantial value in improving decision-making performance when used properly. They are also useful in many aspects of the decision process. Their applicability in situation assessment (especially as probability processing aids) is clear, as is their applicability to generation of options and hypotheses (especially as structuring aids). Structuring aids can also be useful in outcome prediction where no outcome calculator is available, as an unaided estimate of the outcome of a systematically analyzed process is likely to be more accurate than an unaided estimate of the outcome of an incompletely analyzed process. Values of utility assessment techniques are especially useful for interpretation of situational outcomes and resource allocation, but are of little applicability in other parts of the decision process.

4.5.4 Training/Practice Techniques

One special characteristic of military decisions is that many of them occur only in wartime. While it is possible to anticipate the form and structure of C^2 decisions such as strike planning, it is difficult for decision makers to gain experience with the decisions they will have to make in wartime. A special category of decision-aiding techniques, the training/practice

techniques, addresses this problem by providing vehicles by which decision makers can practice by making decisions and obtaining feedback on performance.

Outcome calculators provide this capability to some degree, but by themselves they do not usually provide the real experience of C^2 decision making as a sequence of unfolding events and decisions. Most outcome calculators simply calculate final predicted results from initial conditions. The Strike Outcome Calculator of SRI, however, represents a compromise between outcome calculation and war gaming. After the decision maker establishes an initial campaign plan, the calculator can provide a day-by-day assessment of the status of the campaign, allowing him to change tactics at any time. Thus the decision maker using this aid gets some training in the kinds of compensative decision making that is often required when a sequence of events unfolds and the original plan is found less than optimal.

A more direct method of this sort is the technique of war gaming. The computer can provide a way of simulating a very realistic war game without the extensive facilities needed for more traditional (non-computer-based) war games. The computer may be programmed to choose the strategies of a simulated enemy, or it may be programmed to simulate the interactions of two human decision makers taking opposing sides in the tactical situation. Although this technique was not implemented in any of the ODA projects, it was advocated and incorporated into the design of the comprehensive Air Strike Planner by Analytics (Glenn and Zachary, 1979). This aspect of decision aiding addresses the issue of improving decision quality indirectly by providing the decision maker with realistic experience.

It is interesting to note that there is a reverse way in which training/practice techniques can be used in decision aids -- to have the computer learn the decision-making rules of the human decision maker through observation of his performance in simulated or practice decisions. Thus whereas war gaming uses the computer to train the human, "adaptive" aiding uses a highly skilled human to train the computer. This technique is especially attractive in that it avoids the pitfalls of having decision-aid designers

determine the "proper" or best decision-making strategy. In an adaptive aid, the designer simply models the problem domain and builds an algorithm which allows the computer to learn the human's strategy. One decision-aiding effort has incorporated this approach in an adaptive decision aid (see Weisbrod, Davis, and Freedy (1977)). One problem with this approach is that it is highly dependent on the power of the learning algorithm. If the computer is not able to learn the human's decision rules because its learning algorithm is too simple, then the aid may develop a distinctly suboptimal decision-making procedure. Given the limitations of current state-of-the-art computer learning algorithms, it is unlikely that this adaptive technique will be appropriate for most complex decision-making domains, such as air strike planning, without substantial further research.

4.6 EVALUATION OF DECISION AIDS

A great deal of effort in the ODA program was focussed on the evaluation of decision aids. Projects were devoted to both the development of general evaluation methods and the assessment of specific features of specific aids. Although much useful information was obtained about the aids that were evaluated, probably the most valuable output of these evaluation efforts was an appreciation of the problems that must be confronted by any decision-aid evaluation.

4.6.1 The Criterion Problem

The most basic problem in decision-aid evaluation is the identification of suitable criteria. Evaluation of a decision aid inevitably entails evaluation of a complete decision-making process and that, in turn, necessitates the identification of criteria for decision-making performance. The three general classes of criteria that are usually considered for such evaluations are speed, user acceptance, and decision quality.

Speed of decision making is often singled out as the only objective reliable performance measure that can be obtained. Certainly, faster is better, all other things being equal. And in air strike planning, in particular, speed of decision making is a critical problem. However, it is not generally true in

decision-aid evaluations that all other things are held equal because interesting decision aids tend to have large impacts on user acceptance and decision quality as well as decision speed. Time required for decision making should be recorded whenever possible in decision-aiding experiments for consideration in conjunction with the other performance criteria. It should be recognized, though, that obstacles to accurate determination of decision-making times are often encountered, such as in the evaluation of ASTDA performed by Applied Psychological Services, Inc., in which computer response time variability (the experiment was conducted on a time-sharing system) contaminated data on human decision-making times. Even where fairly clear-cut data can be obtained concerning joint measures of speed and quality, significant interpretation problems can arise. In the ISC evaluations of route planning aids based on Operator-Aided Optimization (OAO) and Iterative Manual Optimization (IMO), it was determined that OAO produced routing decisions that were only slightly superior to those generated by IMO in the same amount of time. But currently projected technological developments in computer hardware capabilities over the next decade could easily enable the speed of an OAO system to be improved by as much as a factor of ten, while comparable increases in IMO speed are unlikely.

Data concerning user acceptance of decision aids are the least expensive and most easily obtained indicators of decision-aid usefulness, so virtually every attempt to evaluate a decision aid collects some subjective data on the attitudes of prospective users toward the usefulness of the aid. Such data is certainly important during the design and redesign phases of aid development since the professed needs and attitudes of the ultimate users should be major determinants of aid specifications. Even when an aid is fully developed, it is appropriate to continue to consult typical users to determine if the aid has indeed addressed a real problem in a realistic fashion. It should be recognized, however, that there are severe limitations to the kinds of useful feedback that can be expected from the operational community because of two potent biasing influences. One bias is the predisposition to reject any aid that requires a significant departure from conventional decision procedures and techniques. A second bias, operating in opposition to the first, is the

tendency to become enamored of any innovation that breaks the tedium of established procedures (i.e., the so-called Hawthorne effect). Since it is impossible to know the magnitude of these two biases in each case, data on user acceptance of decision aids must be treated circumspectly. Appropriate training and indoctrination of aid users can help to minimize the effects of the natural biases, but the magnitude of such mollifying effects is also generally indeterminate. Of course, if the intended users reject the aid over an extended period in an operational setting, it must be concluded that something is seriously wrong; even in that case, however, the culprit may still be the training program rather than the aid design.

Considering the problems in using speed and user acceptance as criteria of decision-making performance, it is always desirable to incorporate some measure of decision quality in the evaluation of a decision aid. Selection of appropriate decision-quality measures, however, is generally a difficult problem since there are many options to choose from.

For aids that offer some prediction of process outcomes, it is desirable to base at least part of the evaluation on a comparison between aid predictions and some independently-obtained "ground truth." However, there are two serious problems which tend to frustrate attempts to evaluate decision-aid predictions against such ground truth:

- Decision aids generally use the best available models for predicting situation outcomes, so how can we find a more reliable method for validation of predictions? In some cases, it might be possible to use large-scale simulations which have been extensively validated but which would be too cumbersome for incorporation in an aid. In other cases, field exercises or war-game simulations might be used to generate the ground truth, but this approach can be very difficult and costly, particularly in the domain of strike operations.

- Because it is generally very difficult to determine the probable outcome or optimal action in a realistically complex situation, decision-making studies often focus on simplified situations for which appropriate outcomes and actions can be determined. This can create a serious problem in extrapolating from the simplified to the realistic situation. ISC recognized this difficulty in

attempting to estimate to what degree real-world strike route planning might be improved based on observed improvements in an experimental setting (Irving et al., 1977).

For aids that primarily offer a vehicle for the implementation of the techniques of decision analysis (e.g., DDI's TACAID, SRI's Decision Structuring Aid, and Grumman's Options Selection Checklist), it has been suggested that the principal criterion for aid evaluation should be coherence of actions and of value and probability assignments (Peterson et al., 1977). The idea is that decision analysis should primarily serve to improve the consistency of decision making. While consistency is certainly an important attribute of a good decision-making system, it must also be recognized that common biases in human judgment can generate internally consistent performance, so reference to some external standard of performance is still desirable.

4.6.2 The Input Problem

An important aspect of decision-aid usefulness is the availability of suitable input data for the aid. Decision analytic models require inputs concerning values, probabilities, and option and event sequences. Outcome models require inputs that describe process parameters and initial conditions. Even if the forms of these models are totally valid, the model outputs can be invalid if the input data is not adequately accurate.

It would be appropriate to systematically examine the sensitivity of decision-aid outputs to perturbations of input data in order to generate specification for necessary input data precision. Of course, it would only be reasonable to undertake this kind of effort after the forms and functions of the aid are firmly established and tentatively accepted for fleet implementation pending verification of system validity. While none of the ODA products have yet achieved a state of development where thorough investigation of input accuracy specifications has been warranted, some studies have examined the feasibility of obtaining the key inputs required by the aids.

ISC examined the capability of humans to generate composite contour diagrams for detection capabilities of a field of radars based on displays of

individual radar capabilities (Irving et al., 1977). They found that the composite diagrams were produced quite accurately and that subsequent route optimization was minimally affected by any imprecision in the human input estimation.

In the evaluation of TACAID, NPRDC required experimental subjects to provide estimates for likelihood ratios for each indicator that was used for Bayesian updating of probabilities of enemy intent (Gettys et al., 1976). Although the likelihood ratios were estimated with reasonable accuracy, it was determined that conditional dependencies between indicators would also have to be estimated in order for the aid to be useful for tactical situation assessment. Other research has established that humans cannot be depended on for accurate estimates of such conditional dependencies, thus posing a requirement for aiding of this estimation process.

It was suggested in an early test plan for ASTDA (Glenn and Zachary, 1978a) that the sensitivity of the outcome calculator on which the aid is based be evaluated with respect to realistic estimates of input error. The idea here was to assume that the outcome model was fully valid and to determine if sufficiently accurate input data could ever be obtained to make the model useful for the strike timing problems which ASTDA was to address. The suggested procedure was to have a group of experts in strike operations construct a set of strike planning scenarios with full specification of all input data required by ASTDA. They would then determine what information would realistically be available to task force personnel who were attempting to estimate the inputs for ASTDA. For example, for the estimation of enemy air defense readiness at candidate strike times, it might be determined that data would be available characterizing enemy air defense readiness on the previous day, effectiveness of attacks on those defenses in the intervening period, and capabilities of the enemy to restore damaged resources. Experimental subjects could then examine the data that characterized each scenario and either (1) make a direct decision about the optimal strike launch time or (2) estimate input data for ASTDA and then run ASTDA to determine the best launch time. The "ground truth" for evaluation of these judgments would be based on the prediction of the ASTDA outcome model using the input data originally generated by the designers of the

scenarios. If the experimental results indicated that the aided decisions were significantly superior to the unaided decisions, then further development of the aid would be warranted. Otherwise, it would be appropriate to develop better techniques for input data estimation before attempting to refine the aid for fleet implementation.

This evaluation of ASTDA based on adequacy of input data accuracies would have been extremely difficult and costly to conduct since verisimilitude of all scenario details, ground truth data, and information made available to experimental subjects would have to be meticulously verified. In fact, this type of evaluation would have been quite inappropriate for ASTDA since such an evaluation just addresses the feasibility of an outcome-model-based aid for strike launch timing and ASTDA was not proposed as an aid just for that application. Rather, ASTDA represented a demonstration of the potential usefulness of a variety of decision-aiding techniques (i.e., outcome modeling, uncertainty band displays, sensitivity analysis) in a representative tactical decision-making situation. Thus the evaluation of ASTDA which was conducted by Applied Psychological Services (Siegel and Madden, 1980) and which focussed on determining the usefulness of particular aiding features in ASTDA was, in fact, more appropriate than the evaluation program described above. However, it is recommended that a program of this type is warranted for any tactical decision aid which incorporates an outcome calculator as a key component and which is seriously considered for operational implementation.

4.6.3 Design Testing and Validation

There are two general issues which any decision-aid evaluation should address:

- Would operational use of the aid produce better decision-making performance than that which obtains without the aid?
- Which features of the aid limit or detract from aid usefulness, and what alternative features would enhance the usefulness of the aid?

The former question addresses the issue of overall aid validity, while the latter question focusses on the issue of how best to improve the design of the aid.

In early phases of aid design, emphasis is appropriately placed on the issue of design improvement, as it was in all evaluations conducted in the ODA program. In evaluating individual features of an aid, it is expedient, if not necessary, to employ a simplified decision situation and a simplified version of the aid in order to achieve the experimental control that is required in order to relate experimental results unequivocally to specific aid features. The research conducted by ISC on route planning aids was very much in this vein of using simplified aid versions and simplified decision problems in order to determine the most effective optimization algorithms, displays of the problem, and allocation of optimization functions between man and machine. Even in the early phases of aid design, however, some efforts to estimate the potential usefulness of the ultimate product must be made to justify continuation of the project, but typically such estimates are made subjectively by potential users or other appropriate experts. An alternative methodology for making more objective estimates of the likely benefits of a decision-aiding concept has recently been proposed in other ONR-sponsored work on decision aiding (Zachary, 1981).

The issue of total aid validity deserves greatest emphasis when the aid design has been firmly established through several iterations of the development and design testing cycle. The important differences between design testing and ultimate aid validation are that final validation requires a more realistic test environment and more tactically relevant performance criteria than does design testing. For design testing, the aid and the problem are generally simplified and the performance criteria are often only indirectly related to the major factors that characterize system performance. But aid validity must ultimately be demonstrated with respect to mission performance of the system in which the aid is embedded. Furthermore, since both the capabilities of decision-making systems (man and machine) and the character of tactical decision problems will always be changing, it is important to recognize that any overall validation of an aid must be recognized to be a continually iterating process much like the early phases of design testing. At any time, new developments in decision-aiding or warfare technology may call for major redesign of an aiding system. In strike route planning, for example, innovations in airborne radars, satellite surveillance, and stealth aircraft would be major new issues

for a system that initially addressed only the routing of conventional aircraft through fields of ship-based and stationary radars.

4.7 MAN-MACHINE INTERFACES FOR DECISION AIDS

Effective interfacing between man and machine was a central concern of all of the decision-aid development projects in the ODA program. Nevertheless, there were very few firm conclusions reached in this work as to the best interface features for military decision aids. This indeterminacy is, in fact, consistent with the general tone of recent reviews of efforts to design man-machine interfaces for computer-based systems (Smith, 1980; Ramsey and Atwood, 1979). There is a plethora of design options and few firm guidelines for selecting between the options to satisfy particular design needs. The problem seems to be that there are many conflicting performance criteria being addressed by decision-aiding systems, so that no one set of interface features can be said to be unequivocally best. Major criteria to be traded off include speed of use, minimization of special training requirements, flexibility of aid for application to diverse problems, adaptability of aid to styles and capabilities of different decision makers, and availability of potentially relevant information. Based on consideration of criteria like these in the context of particular applications, several general principles for interface design appear to be represented in the ODA projects. It can be expected that continuing work in decision aiding will refine these principles and further identify the criterion conditions appropriate for the invocation of each approach.

4.7.1 Provide User Access to all Potentially Relevant Information

In many decision-aid design efforts, the aid is conceived as a versatile tool which should be useful in a wide variety of problems or should support a variety of types of attack on a single type of problem. A user interface to a tactical data base management system like DAISY exemplifies the idea of the versatile aid which provides the decision maker with a mechanism to access all information which might be relevant to the needs of the decision maker. Similarly, the display windowing technique developed in the ODA projects conducted by the Wharton School (whereby a CRT screen is partitioned into rectangular windows which can be treated as independent display devices) provided a

means for simultaneously displaying several separate information domains to a decision maker with complex information needs.

Versatile decision-aid designs for constrained problem domains are offered by EWAR and ASTDA. These aids enable the user to examine and analyze a body of data and models that relate to a particular tactical decision. Both aids embody analytic techniques based on decision analysis and outcome modeling, which, in principle, are capable of determining the optimal decision when all inputs are specified. However, the aid designers seem to have recognized that the operational problems that the aids are intended to address are really more complex than the aid models assume. Thus these aids can provide central frameworks for problem analysis, but it is necessary for the user to have the capability to examine intermediate results and assumptions of the aids to determine how to adjust aid recommendations according to idiosyncrasies of a current problem situation. In this vein, ASTDA provides the user with predictions of outcomes (i.e., force unit attrition) as well as total utility for strike engagements, and the outcome predictions can be broken down by phase of the strike mission. ASTDA also describes both utilities and outcomes in terms of predicted distributions in order to convey information to the user regarding uncertainty and hence risk associated with the predictions. In the domain of task force EMCON planning, EWAR provides a similar variety of interface features to support thorough analyses by the user. These include displays of radar detection capabilities, likely enemy inferences concerning the identity of task force ships, likely results of enemy attacks on the task force, and representations of the trade-off between surveillance capability and information given away by EMCON plans.

4.7.2 Provide the User Only with Basic Necessary Information

The approach to interface design described above assumes that the user has the time and interest to study the fine structure of a problem and assess the impact of idiosyncratic features that cannot be incorporated in a general analysis. A very different design philosophy is appropriate when it is assumed instead that the aid must support quick option selection with a minimum expenditure of effort. Even when the decision may have very important consequences,

it may be possible to identify a few informational inputs that fully characterize the problem and a predetermined decision rule that can be reliably used to determine an adequate selection. This strategy is particularly appropriate for aids that are designed for use by time-constrained individuals like the task force commander and his staff. In fact, it is always appropriate to design an aid to be as efficient as possible in its demands on the user's time and attention. The issue here is whether or not to incorporate features in the aid that are only tangentially related to the basic problem addressed by the aid.

The TACAID represents an example of an aid in which only data necessary for problem resolution is input/output. The tactical problem is defined in minimal detail at a high level. Just a few relevant conditions and action options are considered, and the output is a simple indication of which action should be preferred based on available information. While the triangular output display enables the user to observe the sensitivity of the recommended action to unfolding indicator data, it makes it rather difficult for the user to estimate the probabilities of the conditions (e.g., enemy intent) on which the recommendation is based. It is interesting to note that the triangular output display was generally preferred to the less austere alternative designed by NPRDC in the NPRDC evaluation of TACAID.

4.7.3 Capitalize on Special Human Capabilities

By designing a decision-aid interface to use modalities of information transfer with which humans are particularly adept, it is possible to facilitate effective interface operation. Two modalities worthy of special attention for this purpose are natural language processing and visual interpretation of pictorial displays. Although some efforts have addressed the construction of a natural language interface to a tactical data base (e.g., LADDER, which is described in Sacerdoti and Sagalowicz (1980)), none of the ODA projects employed this interface modality. Many of the ODA efforts did, however, incorporate innovative graphic displays for input and output functions.

Map-based displays of tactical information were employed in the ISC route planning aids, DSA's EWAR, and display concepts recommended by Grumman.

ISC, in fact, proposed the concept of operator-aided optimization as an ideal marriage of the human capability to process complex visual scenes and the computer capability to execute powerful mathematical optimization algorithms. Human aptitude for processing tactical map displays was well-documented in the ISC research which demonstrated that people could:

- generate very good estimates of composite detection rates for a field of radars for which individual radar capabilities were displayed;

- produce very good initial estimates for optimal routes through a field of radars (better than any of the mathematical algorithms designed for this purpose);

- perform route optimization very nearly as well as the best obtainable optimization algorithm when the computer produces the radar detection display and provides feedback on the utility of candidate routes.

Another type of display innovation developed in the ODA program is the delta-biased uncertainty band, which was used in ASTDA to represent statistical distributions in an efficient manner (Glenn and Zachary, 1978b). This display technique consists of displaying the mean of a distribution as a point and presenting a two-standard-deviation band around the mean, with the band being offset so as to characterize the skewness of the distribution. It was proposed that this type of display would enable aid users to compare distributions visually and to assess the likelihood that observed values were consistent with predicted distributions.

4.7.4 Help Overcome Human Memory and Attention Limitations

Interface features can be used to alleviate problems associated with human cognitive limitations. Inabilities to reliably remember procedural steps or factors that should be considered in problem analysis can be countered by use of rigid operation sequences, prompts, and checklists. The SRI Strike Outcome Calculator used a rigid operation sequence to ensure that all necessary data was provided. The SRI Decision Structuring Aid used prompts to help the decision maker recognize important decision factors. The checklist is similarly

used in Grumman's Options Selection Checklist to ensure that all appropriate factors are included in the analysis.

4.7.5 Help the User Understand How the Aid Works

A serious obstacle to effective use and user acceptance of aids that have complicated structures or that perform esoteric types of processing is for the user to learn and remember how the aid operates and what it can and cannot do. Potentially useful features are of no value if the user forgets that they are available or cannot invoke them. A major innovation to address this problem is the influence diagram which SRI developed as an analysis tool for the Decision Structuring Aid. The influence diagram is a type of network display which illustrates how decision problem factors impact one another. SRI and DSA collaborated on the incorporation of the influence diagram in EWAR to portray the structure of the EMCON planning problem and the relationships of aid components to aspects of the problem structure. This feature constitutes a sort of built-in training device for the aid, serving to train the user both in general understanding of the problem domain as well as in specific organization and capabilities of the aid.

Other tools which help the user understand the operation of a process model are sensitivity analysis and intra-process analysis. Sensitivity analysis permits the user to observe how model outputs vary in response to variations in one or more model inputs. Intra-process analysis enables the user to examine how intermediate outcomes unfold in a time-extensive process. Both of these analysis tools are offered in ASTDA to enable the user to understand the operation of the engagement model on which ASTDA is based. Sensitivity analysis is incorporated as an option in the Decision Structuring Aid as a tool for examining the impact of uncertainties in probability and utility assignments on recommended choices. Sensitivity analysis is also a component of EWAR, which aids in the interpretation of model interactions. By enabling the user to assume selective control over models, sensitivity analysis can help him understand the behavior of the model and assess whether or not it agrees with the user's intuitions about the process of concern.

5. ISSUES FOR FURTHER RESEARCH

The ODA program has performed an important function in developing decision-aiding techniques and evaluation methods that have a wide range of applications, particularly in the area of strike planning. In the course of reviewing efforts conducted in the ODA program, it has become clear that decision-aiding technology holds great promise for the operational community, but that several issues require further intensive investigation before many of the potential benefits of decision aiding can be realized. Some of these issues are identified and discussed in this section.

5.1 MAKING USE OF FLEET KNOWLEDGE ABOUT FLEET PROBLEMS

It is important to develop effective techniques to use the knowledge of experienced operational decision makers regarding what problems are in need of aiding and what types of aids are best suited for those applications. In the domain of air strike planning, in particular, it is clear that there is not general agreement as to what decision-making problems cause the greatest difficulty. For example, some strike planners interviewed by Analytics have indicated that route planning is a trivial problem and that weaponeering is very difficult, while others indicated that the reverse was true. These differences probably result from different problem environments being considered by different individuals (e.g., environments in which routing and weaponeering options may be extremely rich or extremely limited). It is appropriate to develop structured methods to elicit information on aiding needs from individuals that somehow account for possible individual biases and also to design procedures to aggregate inputs from many individuals to construct a representative picture of total fleet needs.

5.2 AIDING IN GOAL CONSIDERATION

Most decision-aiding efforts assume that decision-making goals are best relegated to subjective consideration by the decision maker. Except for helping in the trade-off between distinct outcome criteria which are presumably related to conflicting goals, the aids produced in the ODA program did not address the problem of analyzing a complex system of goals as an aid to decision analysis. In the case of strike warfare, there is definitely a complex network of goals with definable interrelations (e.g., arrive at target undetected, deliver ordnance on target, inflict target destruction, avoid own force losses, restrict enemy options, etc.). It is desirable to aid decision makers in systematic representation and analysis of such goal systems in order to improve decision-making consistency and sensitivity of intermediate decisions to ultimate goals. One interesting effort in this area has been reported by Pearl et al. (1980). Also, much of the recent research in artificial intelligence on goal-directed computation (see, for example, Wilensky (1981)) may be relevant here as well. It appears that further research efforts in these areas with application to specific military goal domains are warranted.

5.3 MATCHING AIDING TECHNIQUES TO PROBLEM CHARACTERISTICS

Methods are needed to determine which decision-aiding techniques are most suitable for particular decision problem applications. Under what conditions are decision analytic models or outcome calculators appropriate? What level of granularity should be adopted for the problem representation? What kinds of auxiliary analysis techniques should be offered to the decision maker? A number of research efforts have addressed the problem of matching decision-aiding techniques to problem characteristics, both within the ODA program (Pugh, 1976; Payne et al., 1975; Peterson et al., 1977) and in other ONR-sponsored efforts (Miller et al., 1980 Brown and Ulvila, 1977; Zachary, 1980), and coarse correspondences between techniques and problem features have been identified. Further refinement of these matching procedures should be pursued, probably based on a distillation of the results of empirical evaluations of decision-aid techniques.

5.4 AIDING IN GENERATION OF DECISION OPTIONS

Virtually all decision-aiding research conducted in or out of the ODA program has assumed that the decision maker starts the decision process with knowledge of his goals and viable action options and that the purpose of a decision aid is to assist in the judicious selection of an action which is most suitable for the achievement of the goals. It has already been noted that more research is needed on goal consideration. Similarly, much more research is warranted to develop techniques to aid in the generation or identification of decision options. In any problem situation like strike planning, for example, the feasible option space is restricted and structured by resource and environmental factors (e.g., available aircraft, acceptable tactics). Promising techniques to aid in the generation of options include the artificial intelligence techniques of rule-based inference (which has been used in an experimental aid called KNOBS developed for Air Force strike planning -- Engleman et al. (1979)) and problem-solving algorithms. Innovations in graphic display technology may also serve to aid in option generation, as ISC's route planning aids may be considered as a sort of option generation tool.

5.5 VALIDATION OF OUTCOME CALCULATORS

Since a large class of decision aids is based on the use of outcome calculators, an important issue for decision-aiding research is how to ensure that an outcome calculator is adequately valid for a particular decision-aiding application. While empirical validation of outcome calculators is a major subject of investigation for operations research and military science, the determination of minimal validity specifications for a model to be used in a particular decision-aiding application is a special problem for decision-aiding researchers. When the outcome calculator has minimal internal complexity (such as in the case of SRI's Strike Outcome Calculator), this problem reduces to that of defining accuracy specifications for inputs to the model as discussed in Subsection 4.6.2. When the outcome model is more complex (as in the case of ASTDA), it is necessary to specify accuracy limits for both the inputs and the outputs of the model.

5.6 EVALUATION OF PREFERENCE-MODELING TECHNIQUES

A majority of the decision aids developed in the ODA program employ value models for the preferences of decision makers as a key component. Grumman's Options Selection Checklist, in particular, was little more than a mechanism for using a multi-attribute utility model to quantify a decision maker's preferences. Research is required to determine the reliability of such preference-modeling techniques in tactical military applications. How consistently can value assignments be made? Are decisions made with the help of such preference-modeling tools in better agreement with high-level goals than unaided decisions? What analytic forms are best for multi-criteria utility functions? Hard evidence of the validity of preference-modeling methods will be required to overcome the resistance to these techniques that has been expressed by some operational decision makers (e.g., in the evaluation of ASTDA reported by Siegel and Madden (1980)).

5.7 USEFULNESS OF UNCERTAINTY INFORMATION

Two of the ODA projects developed techniques for the display of information characterizing the degree of uncertainty associated with outcomes predicted by a model. NPRDC designed an alternative to the TACAID display in order to present probabilities associated with distinct decision results. Although the NPRDC evaluation of TACAID indicated that subjects preferred the DDI version of TACAID, which did not present uncertainty information, that finding may result from a variety of factors which do not reflect the usefulness of uncertainty information (e.g., the NPRDC display of outcome probabilities may have obscured the display of recommended action). ASTDA offered a different type of display of uncertainty information to convey salient features of probability distributions to aid users. It was suggested that the uncertainty bands of ASTDA would help the user of the aid to assess the accuracy of aid predictions based on actual observed outcomes. This hypothesis was not, however, addressed by the experimental evaluation of ASTDA. Further investigation into the usefulness of such displays of uncertainty is warranted to determine the conditions under which such information is and is not beneficial and the best ways for displaying the information when it is relevant. If the information is extraneous to the decision process, it should be excluded from the aid design to simplify the

interface and the processing algorithms. But care should be taken to examine all conditions under which uncertainty data might be useful to the decision maker before eliminating information which may occasionally be very important.

5.8 ALLOCATING FUNCTIONS AND DESIGNING INTERFACES

Central issues in the design of decision-aiding systems are how best to allocate information processing functions between the human and the machine and how to design effective interfaces for the transfer of information between man and machine. Although these issues have long been fundamental concerns of human factors engineers in the design of man-machine systems, special problems warranting new research efforts are encountered in the design of decision-aiding systems. Decision aids often call for the sharing of intelligent functions between man and machine, but reliable characterizations of relative human and machine capabilities to accomplish such functions are difficult to obtain. Technology is advancing rapidly in areas of artificial intelligence and display and control hardware, making it inadvisable to base design decisions on empirical performance data of candidate-aiding components in many cases. Recent surveys of work concerning these problems is useful (e.g., Ramsey and Atwood, 1979; Smith, 1980), but new comprehensive methods for designing decision-aiding systems are needed. Initial observations from one such effort are reported by Zachary et al. (1981).

5.9 PRIORITIZING CANDIDATES FOR DECISION-AID DEVELOPMENT

Resources for development of new systems are always limited, and problems that could potentially be alleviated by decision-aiding tools will generally compete for those resources. Military program managers must consequently make decisions about which decision-aid developments are likely to produce the most cost-effective improvements in overall system performance. In order to support such resource allocation decisions, it is desirable to develop an objective reliable methodology for assessing and comparing the probable benefits and costs for alternative candidate decision-aid concepts. Such a methodology must address the formidable problem of estimating the overall impact on system performance of coarsely-defined aiding concepts and the various costs

associated with development, implementation, and maintenance of operational aiding systems. One ONR-sponsored effort to develop this type of methodology has been reported by Zachary (1981).

5.10 DESIGNING AIDS THAT LEARN

Many decision aids are designed for frequent use to address chronic or recurring problems. At the same time, decision aids typically require the user to provide input data estimates for which the user has imprecise information and no systematic estimation procedures. By observing the actual outcomes that obtain in one iteration of a decision process, it should be possible for the decision maker to adjust input estimates so as to improve their accuracy for the next decision of the same type. However, since humans are often very poor at using feedback to make adjustments to initial estimates and since efficient mathematical algorithms are available for this purpose (e.g., Kalman filtering and Bayesian updating), it is appropriate to build such a learning feature into a decision aid. In strike planning, for example, cyclic strikes are required in many situations so that a strike planning aid could benefit from the capability to adjust estimates of enemy force readiness and measures of engagement capabilities for own and enemy forces on the basis of results from earlier strikes.

6. SUMMARY AND CONCLUSIONS

At its inception, the focus of the ODA program was on the decision-making problems of the task force commander and potential aids to alleviate those problems and improve decision-making performance. Several aids to address high-level problems like those of the flag command were developed. However, the focus of the program gradually changed until, near the end of the program, most of the ODA efforts were concerned with lower-level tactical problems, particularly the various aspects of strike planning. With the redirection of emphasis toward the detailed considerations of lower-level problems, outcome calculators were introduced into the decision analytic formulations that had been developed for high-level decisions. Decision analytic aids conceived early in the program include DDI's TACAID, Grumman's Options Selection Checklist, and SRI's Decision Structuring Aid. SRI's Strike Outcome Calculator is an aid based on a collection of models that is suitable for strike campaign planning at an intermediate organizational level (e.g., the Strike Planning Board). Other outcome model aids are DSA's EWAR for task force EMCON planning, ISC's Air Strike Planning System, and Analytics' Air Strike Timing Decision Aid. A number of general-purpose informational supports for decision makers were also developed in the course of the program. Many experimental evaluations of decision aids were conducted in the ODA program, including evaluations of three aids by independent contractors.

It is very difficult to assess the degree of technological progress in decision aiding for which the ODA program is responsible, but it is clear that its contribution has been very great. Before the ODA program there was no coherent discipline of decision-aiding research, and now there is a very active one. ODA succeeded in its expressed objective of bringing about the coordinated interaction of researchers from the separate disciplines of operations research, decision analysis, computer science, and human factors engineering

to develop tools to support military decision makers. A broad variety of decision-aiding techniques were developed which illustrate the diversity of techniques and applications which decision-aiding research can offer. Important lessons were also provided as to how to design and evaluate such decision-aiding systems.

Many of the ODA efforts focussed on providing support to air strike planners. Applications of the decision analytic techniques of value modeling and probability assessment to strike planning problems were investigated. A variety of strike force routing aids were constructed and evaluated in order to determine appropriate roles for human judgment, innovative graphic display techniques, and mathematical optimization algorithms. The problem of selecting an advantageous strike launch time was addressed by one decision-aiding system. Another aid provided a tool for defining and estimating the likely results of time-extensive multi-strike campaigns. Several concepts were proposed for comprehensive air strike planning systems. Informal consultations with Navy officers who have been involved in strike planning make it clear that some type of decision aiding is needed in this area, but it is not clear whether any of the aids developed to date really address the needs. In the future, it is important to make concerted efforts to determine the requirements and constraints for strike planning assistance before any aid is developed for operational implementation. Determination of these requirements and constraints must depend both on structured solicitation of fleet inputs and continuing experimental research in decision-aiding technology.

A number of important questions for further research have been identified here. These include:

- How can we make more effective use of the knowledge of experienced decision makers to determine what problems warrant aiding and what types of aids are best?

- How can decision-making processes be made more sensitive to appropriate goals?

- What is an effective methodology for matching decision-aiding techniques to specific problem characteristics in order to aid the process of decision-aid design?

- How can we provide help to decision makers in generation or identification of viable decision options?

- For decision aids that incorporate outcome calculators, how can we develop procedures to determine the degree of validity which is required of the outcome calculator in order for the aid to be useful, and how can the appropriate validation be accomplished?

- What are reliable techniques for evaluation of preference-modeling techniques?

- Under what, if any, conditions is it beneficial to offer information to a decision maker concerning the uncertainty associated with model-based predictions of outcomes?

- What are useful general principles for allocation of cognitive responsibilities between man and machine in decision-making systems and for design of man-machine interface in such distributed-intelligence systems?

- How can we develop useful techniques for prioritizing candidate decision-aiding concepts for an applications area?

- How can we develop model-based aids to adjust model parameters based on feedback from an unfolding sequence of decision processes and, thus, to iteratively improve aid performance?

Research in some of these areas is currently in progress, and research in the remaining areas is warranted immediately. It is apparent that the high level of interest and activity in decision-aiding research that was begun by the ODA program will be continued with productive results long after the termination of that program.

AD-A121 939 DESIGN OF AIR STRIKE PLANNING AIDS: LESSONS LEARNED FROM THE ONR OPERATIO..(U) ANALYTICS INC WILLOW GROVE PA F A GLENN ET AL. 15 MAY 82 1454-B N00014-79-C-0656 2/2

UNCLASSIFIED F/G 5/1 NL

END

FILMED

DTIC

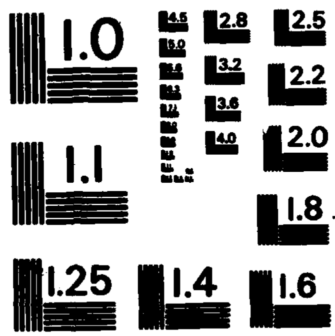

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

REFERENCES

1. Barclay, S., C.R. Peterson, L.S. Randall, and M.L. Donnell. *Decision analysis as an element in an operational decision aiding system: Phase V* (Tech. Rep. PT 78-29-6). Decisions and Designs, Inc., April 1979. (AD A068 339)

2. Brown, R.V., C.M. Hoblitzell, C.R. Peterson, and J.W. Ulvila. *Decision analysis as an element in an operational decision aiding system* (Tech. Rep. 74-2). Decisions and Designs, Inc., September 1974. (AD A001 110)

3. Brown, R.V., C.R. Peterson, W.H. Shawcross, and J.W. Ulvila. *Decision analysis as an element in an operational decision aiding system (Phase II)* (Tech. Rep. 75-13). Decisions and Designs, Inc., November 1975. (AD A018 109)

4. Brown, R.V., and J.W. Ulvila. *Selecting analytic approaches for decision situations* (3 volumes) (Tech. Rep. TR77-7-25). Decisions and Designs, Inc., 1977.

5. Buneman, P.O., and H.L. Morgan. *Implementing alerting techniques in data-base systems* (Tech. Rep. 77-03-04). University of Pennsylvania, 1977.

6. CTEC, Inc. *Office of Naval Research operational decision aids (ODA) program: ODA data base*. CTEC, Inc. (Tech. Rep. 26056). February 1976. (AD B010 601)

7. CTEC, Inc. *Information support for operational decision aids*. (Tech. Rep. 57283). CTEC, Inc., May 1977. (AD A041 223)

8. Densmore, J.E., R.M. Kerchner, D.F. Noble, G.E. Pugh, and P.G. Tomlinson. *An emissions control decision aid* (Vol. I) (Tech. Rep. DSA-66). Decision-Science Applications, Inc., July 1978. (AD A060 072)

9. Elam, J.J. *Model management systems: A framework for development* (Tech. Rep. 79-02-04). University of Pennsylvania, February 1979. (AD A067 246)

10. Engleman, C., S. Bergrand, and M. Bischoff. "KNOBS: An experimental knowledge based tactical air mission planning system and a rule based aircraft identification simulation facility." In *Proceedings of the Sixth International Joint Conference on Artificial Intelligence,* Tokyo, Japan, 1979. Vol. I, p. 247.

REFERENCES (continued)

11. Epstein, S., M. Strieb, R. Goldman, F. Glenn, and R. Wherry. *Operational decision aids: The application of nomography and uncertainty analysis to decision-aiding systems* (Tech. Rep. 1218-A). Analytics, November 1977. (AD A049 545)

12. Garnero, R.S., J.C. Bobick, and D. Ayers. *Strike outcome calculator (SOC) -- Description and operating instructions* (Tech. Rep. NWRC TR-15). SRI International, March 1978a. (AD A061 363)

13. Garnero, R.S., J.V. Rowney, and J. Ketchel. *Evolution and preliminary tests of the strike outcome calculator* (SOC) (Tech. Rep. NWRC TR-16). SRI International, October 1978b. (AD A061 364)

14. Gettys, C.F., M.C. Moy, and M.W. O'Bar. *Significance of risk in Navy tactical decision making: An empirical investigation* (Tech. Rep. NPRDC TR 77-8). Navy Personnel Research and Development Center, December 1976.

15. Glenn, F. *Experimental strike timing decision problems* (Tech. Rep. 1344-B). Analytics, 1978.

16. Glenn, F., and J. Bennett. *Decision aiding concepts for air strike planning* (Tech. Rep. 1454-A). Analytics, 1980.

17. Glenn, F., and W. Zachary. *Test plan for the Analytics strike timing decision aid* (Tech. Rep. 1212-B). Analytics, 1978a.

18. Glenn, F., and W. Zachary. *ASTDA user's guide* (Tech. Rep. 1344-A). Analytics, 1978b.

19. Glenn, F., and W. Zachary. *Integration of decision aids for strike campaign planning* (Tech. Rep. 1344-C). Analytics, May 1979. (AD A069 751)

20. Hurst, E.G. (Ed.). *Description of the Wharton/ODA system* (WP 77-1-02). University of Pennsylvania, November 1977. (AD A046 404)

21. Hurst, E.G., H.L. Morgan, and D.N. Ness. *Decision aiding information system (DAISY) user's guide* (Working Paper 75-01-02). Wharton School, University of Pennsylvania, January 1975a.

22. Hurst, E.G., H.L. Morgan, D.N. Ness, and R.J. Zowader. *DAISY/APL interface user's memo* (Working Paper 75-01-04). Wharton School, University of Pennsylvania, January 1975b.

23. Hurst, E.G., H.L. Morgan, and D.N. Ness. *DAISY: A decision-aiding information system* (Working Paper 75-01-05). Wharton School, University of Pennsylvania, January 1975c.

REFERENCES (continued)

24. Irving, G.W. *Continuous subjective function (CSF) pilot study* (Tech. Memo). Integrated Sciences Corporation, September 1975.

25. Irving, G.W., J.J. Horinek, A.J. Himmel, P.Y. Chan, and D.H. Walsh. *Experimental investigation of sketch model applications to tactical decision aiding* (Tech. Rep. 215-3). Integrated Sciences Corporation, January 1977.

26. Irving, G.W., J.J. Horinek, D.H. Walsh, and P.Y. Chan. *ODA Pilot Study II: Selection of an interactive graphics control device for continuous subjective functions applications* (Tech. Rep. 215-2). Integrated Sciences Corporation, April 1976.

27. Kalenty, C.R., and W.L. Lockwood. *Experimental evaluation of an options matrix as an operational decision aid* (Tech. Rep. CSS 76-3). Grumman Aerospace Corporation, March 1976.

28. Kalenty, C.R., W.L. Lockwood, and V.M. Vissering, Jr. *Experimental validation of an options selection matrix and investigation of other display formats as operational decision aids* (Tech. Rep. CSS 77-1). Grumman Aerospace Corporation, February 1977.

29. Lucas, G.L., and J.A. Ruff. *An investigation of operational decision aids* (Tech. Rep. 312). System Planning Corporation, July 1977. (AD 047 147)

30. Madden, E.G., and A.I. Siegel. *Evaluations of operational decision aids: II. The emissions control aid* (Tech. Rep.). Applied Psychological Services, Inc., April 1980.

31. Martin, W.H., Jr. *Measuring the performance of operational decision aids* (Tech. Rep. 1161-B). Analytics, April 1976. (AD A024 795)

32. Merkhofer, M.W., and E.B. Leaf. *A computer-aided decision structuring process -- final report* (Tech. Rep. 1513). SRI International, June 1981.

33. Merkhofer, M.W., A.C. Miller, R.A. Howard, and S.N. Tani. *A preliminary characterization of a decision structuring process for the task force commander and his staff* (Tech. Rep. 4030). Stanford Research Institute, December 1975. (AD A019 302)

34. Merkhofer, M.W., A.C. Miller, III, B.E. Robinson, and R.J. Korsan. *Decision structuring aid: Characterization and preliminary implementation.* SRI International, September 1977.

REFERENCES (continued)

35. Merkhofer, M.W., B.E. Robinson, and R.J. Korsan. *A computer-aided decision structuring process* (Tech. Rep.). SRI International, June 1979. (AD A072 146)

36. Miller, A., P.A. Morris, R.D. Smallwood, and R.S. Gibbons. *Analytic procedures for designing and evaluating decision aids* (Tech. Rep.). Applied Decision Analysis, Inc., 1980.

37. Miller, L.W., R.J. Kaplan, and W. Edwards. "JUDGE: A value-judgment-based tactical command system." *Organizational Behavior and Human Performance*, 1967, 2, pp. 329-374.

38. Mitchell, R.M., Jr. *A model interface for the decision aiding information system* (Tech. Rep. 76-09-11). University of Pennsylvania, September 1976.

39. Newell, A., J. Shaw, and H. Simon. "GPS: A program that simulates human thought." In E. Feigenbaum and J. Feldman (Eds.), *Computers and Thought*. New York: McGraw-Hill, 1963.

40. Noble, D.F. *Development and evaluation of an emissions control decision aid* (Tech. Rep.). Decision-Science Applications, Inc., July 1980.

41. Noble, D.F., and G.E. Pugh. *A prototype interface to adapt decision aids to user scenario assumptions* (Tech. Rep. DSA-335). Decision-Science Applications, Inc., July 1981.

42. Oppenheim, D., J.S. Ribeiro, and E.G. Hurst. *Description of the Wharton/ODA system* (Tech. Rep. 79-01-02). University of Pennsylvania, December 1979.

43. Payne, J.R., T.J. Braunstein, J.M. Ketchel, and M.C. Pease. *A brief survey of potential decision aids for the task force commander and his staff* (NWRC-RM-84). Stanford Research Institute, August 1975.

44. Payne, J.R., A.C. Miller, and J.V. Rowney. *The naval task force decision environment*. Stanford Research Institute, September 1974. (AD A001 873)

45. Payne, J.R., and J.V. Rowney. *ONRODA warfare scenario* (NWRC-RM-83). Stanford Research Institute, June 1975.

46. Pearl, J., A. Leal, and J. Saleh. *GODDESS: A goal-directed decision structuring system* (Tech. Rep. UCLA-ENG-CSL-8034). UCLA, June 1980.

47. Peterson, C.R., L.D. Phillips, L.S. Randall, and W.H. Shawcross. *Decision analysis as an element in an operational decision aiding system (Phase IV)* (Tech. Rep. 77-4-6). Decisions and Designs, Inc., April 1977. (AD A045 083)

REFERENCES (continued)

48. Peterson, C.R., L.S. Randall, W.H. Shawcross, and J.W. Ulvila. *Decision analysis as an element in an operational decision aiding system (Phase III)* (Tech. Rep. 76-11). Decisions and Designs, Inc., October 1976. (AD A034 037)

49. Pugh, G. *Mathematical decision aids for the task force commander and his staff* (Tech. Rep. 593W-01-CR). General Research Corporation, 1976.

50. Pugh, G.E., R.M. Kerchner, P.G. Tomlinson, and D.F. Noble. *An emissions control decision aiding system for fleet air defense: An R&D status report* (Tech. Rep. 912-01-CR). General Research Corporation, May 1977.

51. Raiffa, H. *Decision analysis: Introductory lectures on choice under uncertainty*. Reading, MA: Addison Wesley, 1968.

52. Ramsey, H.R., and M.E. Atwood. *Human factors in computer systems: A review of the literature* (Tech. Rep. SAI-79-111-DEN). Science Applications, Inc., September 1979. (AD A075 679)

53. Ribeiro, J.S. *DBALERT: An alerting system for WAND* (Tech. Rep. 78-12-02). University of Pennsylvania, December 1978. (AD A067 247)

54. Rowney, J.V. *Amphibious warfare scenario* (NWRD-RM-86). Stanford Research Institute, October 1975.

55. Rowney, J.V., R.S. Garnero, and J.C. Bobick. *Augmentation of the naval task force decision aiding system: The outcome calculator*. Stanford Research Institute, April 1977. (AD A039 917)

56. Sacerdoti, E.D., and D. Sagalowicz. *A LADDER user's guide (revised)* (Tech. Note 163R). SRI International, March 1980.

57. Schechterman, M.D., and D.H. Walsh. *Comparison of operator aided optimization with iterative manual optimization in a simulated tactical decision aiding task* (Tech. Rep. 215-6). Integrated Sciences Corporation, July 1980.

58. Siegel, A.I., and E.G. Madden. *Evaluations of operational decision aids I. The strike timing aid* (Tech. Rep.). Applied Psychological Services, Inc., January 1980. (AD A080 835)

59. Sinaiko, H.W. *Operational decision aids: A program of applied research for Naval command and control systems* (Tech. Rep. TR-5). Smithsonian Institution, June 1977. (AD A042 091)

60. Slovic, P., B. Fischhoff, and S. Lichtenstein. "Behavioral decision theory." *Annual Review of Psychology*, 1979, 28, pp. 1-39.

61. Smith, S.L. *Requirements definition and design guidelines for man-machine interface in C^3 system acquisition* (Tech. Rep. M80-10). MITRE Corporation, June 1980. (AD A087 258)

62. Spector, B., R.E. Hayes, and M.J. Crain. *The impact of computer-based decision aids on organizational structure in the task force staff* (Tech. Rep. CAC 210). CACI, Inc., September 1976. (AD A031 654)

63. Walsh, D.H., G.J. Rebane, L.R. Levi, and R. Tash. *Concepts for highly interactive air strike planning systems (ASPS)* (Tech. Rep. 330-2). Integrated Sciences Corporation, May 1981.

64. Walsh, D.H., and M.D. Schechterman. *Experimental investigation of the usefulness of operator aided optimization in a simulated tactical decision aiding task* (Tech. Rep. 215-4). Integrated Sciences Corporation, January 1978. (AD A053 336)

65. Walsh, D.H., and M.D. Schechterman. *An investigation of selected alternative decision aids* (Tech. Rep. 215-5). Integrated Sciences Corporation, April 1979. (AD A070 524)

66. Weisbrod, R.L., K. Davis, and A. Freedy. "Adaptive utility assessment in dynamic processes: An experimental evaluation of decision aiding." *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-7, No. 5, 1977, pp. 377-383.

67. Wilensky, R. "A model for planning in everyday situations." In *Proceedings of the Third Annual Conference of the Cognitive Science Society*, Berkeley, CA, 1981. pp. 11-16.

68. Zachary, W.W. *Decision aids for naval air ASW* (Tech. Rep. 1366A). Analytics, March 1980.

69. Zachary, W.W. *Cost-benefit assessment of candidate decision aids for naval air ASW* (Tech. Rep. 1366C). Analytics, April 1981.

70. Zachary, W.W., F. Glenn, and J. Hopson. "Intelligent man talks to intelligent machine: Implications of distributed-intelligence systems for the design of man-machine interface." In *Proceedings of the Seventh Conference of the Canadian Man-Computer Communications Society*, Waterloo, Ontario, 1981. pp. 99-104.

DISTRIBUTION LIST

Director, Engineering Psychology
Programs (Code 455)
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217 (5 cys)

Defense Technical Information Center
Cameron Station, Bldg. 5
Alexandria, VA 22314 (12 cys)

CDR Paul Chatelier
Office of the Deputy Under Secretary
of Defense
OUSDRE (E&LS)
Pentagon, Room 3D129
Washington, D.C. 20301

Dr. Stuart Starr
Office of the Under Secretary of Defense
(C3I)
Pentagon, Room 3C200
Washington, D.C. 20301

Dr. Craig Fields
Director, Cybernetics Technology Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Office of the Chief of Naval
Operations, OP987H
Personnel Logistics Plans
Washington, D.C. 20350

Dr. A.L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
Code RD-1
Washington, D.C. 20380

Commanding Officer
ONR Eastern/Central Regional Office
ATTN: Dr. J. Lester
Bldg. 114, Section D
666 Summer Street
Boston, MA 02210

Commanding Officer
ONR Branch Office
ATTN: Dr. Charles Davis
536 South Clark Street
Chicago, IL 60605

Commanding Officer
ONR Western Regional Office
ATTN: Dr. E. Gloye
1030 East Green Street
Pasadena, CA 91106

Commanding Officer
ONR Western Regional Office
ATTN: Mr. R. Lawson
1030 East Green Street
Pasadena, CA 91106

Tactical Development & Evaluation
Support
Code 230
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Operations Research Program
Code 434
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Statistics and Probability Program
Code 436
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Information Systems Program
Code 437
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

CDR Kent Hull
Code 430B
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

CDR P.M. Curran
Code 270
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Dr. James McGrath
Code 311
Navy Personnel Research and
Development Center
San Diego, CA 92152

Management Support Department
Code 210
Navy Personnel Research and
Development Center
San Diego, CA 92152

Naval Electronics Systems Command
Human Factors Engineering Branch
Code 4701
Washington, D.C. 20360

Director
Naval Research Laboratory
Technical Information Division
Code 2627
Washington, D.C. 20375 (6 cys)

Commander, Naval Electronics
Systems Command
Command and Control Division
Code 530
Washington, D.C. 20360

Dr. John Silva
Head, Human Factors Division
Naval Ocean Systems Center
San Diego, CA 92152

Dr. Jesse Orlansky
Institute for Defense Analyses
400 Army-Navy Drive
Arlington, VA 22202

Human Factors Department
Code N215
Naval Training Equipment Center
Orlando, FL 32813

Dr. Gary Poock
Operations Research Department
Naval Postgraduate School
Monterey, CA 92940

Dr. Joseph Zeidner
Technical Director
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. H.W. Sinaiko
Smithsonian Institution
801 N. Pitt Street
Alexandria, VA 22314

Office of the Chief of Naval
Operations OP942
Pentagon
Washington, D.C. 20350

Office of the Chief of Naval
Operations OP987C
R&D Plans Division
Washington, D.C. 20350

Commander
Naval Electronics Systems Command
C3 Project Office
PME 10801
Washington, D.C. 20360

Dr. Julie Hopson
Human Factors Engineering Division
Code 604
Naval Air Development Center
Warminster, PA 18974

Dr. M.L. Metersky
Naval Air Development Center
Code 5424
Warminster, PA 18974

Dr. Edgar M. Johnson, Director
Ogranizations & Systems Research
Laboratory
U.S. Army Research Laboratory
5001 Eisenhower Avenue
Alexandria, VA 22333

Mr. Victor Monteleon
Naval Ocean Systems Center
San Diego, CA 91252

Commander, Naval Electronics
Systems Command
ELEX-03
Washington, D.C. 20360

CDR Richard Schlaff
NIPSSA
Hoffman Bldg. #1
2461 Eisenhower Avenue
Alexandria, VA 22331

Dr. Chantee Lewis
Management Department
Naval War College
Newport, RI 02840

Dr. John Shore
Naval Research Laboratory
Code 5403
Communications Sciences Division
Washington, D.C. 20375

Dr. Meredith Crawford
American Psychological Association
Office of Educational Affairs
1200 17th Street, N.W.
Washington, D.C. 20036

Dr. William Dejka
ACCAT
Naval Ocean Systems Center
San Diego, CA 92152

Mr. Merlin Malehorn
Office of the Chief of Naval
Operations, (Op 102)
Washington, D.C. 20350

Dr. S.D. Epstein
Analytics
2500 Maryland Road
Willow Grove, PA 19090

Dr. Amos Freedy
Perceptronics, Inc.
6271 Variel Avenue
Woodland Hills, CA 91364

Dr. G. Hurst
University of Pennsylvania
Wharton School
Philadelphia, PA 19174

Dr. Miley Merkhofer
Stanford Research Institute
Decision Analysis Group
Menlo Park, CA 94025

Mr. George Pugh
Decision Science Applications, Inc.
1500 Wilson Boulevard
Arlington, VA 22209

Dr. Arthur Siegel
Applied Psychological Services
Science Center
404 E. Lancaster Street
Wayne, PA 19087

LCDR J.A. Sears
Department of MIS
College of Business Administration
University of Arizona
Tucson, AZ 85721

I.R. Mirman
Asst for Special Projects
HQ AFSC-DL
Andrews AFB, MD 20334

Mr. Joseph Wohl
MITRE Corporation
Box 208
Bedford, MA 01730

Dr. Kenneth Gardner
Applied Psychology Unit
Admiralty Marine Technology
Establishment
Teddington, Middlesex TW11 OLN
ENGLAND

Mr. Tim Gilbert
MITRE Corporation
1820 Dolly Madison Blvd.
McLean, VA 22102

Mr. Leslie Innes
Defense & Civil Institute of
Environmental Medicine
P.O. Box 2000
Downsview, Ontario M3M 3B9
Canada

Dr. Rex Brown
Decision Science Consortium
Suite 421
7700 Leesburg Pike
Falls Church, VA 22043

Dr. A.C. Miller, III
Applied Decision Analysis
3000 Sand Hill Road
Menlo Park, CA 94025